Artif Intell Rev (2007) 27:257–271
DOI 10.1007/s10462-008-9080-7

An evolutionary approach to deception in multi-agent systems

Pilib Ó Broin · Colm O'Riordan

Published online: 3 September 2008

Abstract Understanding issues of trust and deception are key to designing robust, reliable multi-agent systems. This paper builds on previous work which examined the use of auctions as a model for exploring the concept of deception in such systems. We have previously described two forms of deceptive behaviour which can occur in a simulated repeated English auction. The first of these types of deception involves sniping or late bidding, which not only allows an agent to conceal its true valuation for an item, but also potentially allows it to win an item for which it may not possess the highest valuation. The second deceptive strategy involves the placing of false bids which are designed to reduce an opponent's potential profit. In this work we examine the potential shortcomings of those two strategies and investigate whether or not their individual strengths can be combined to produce a successful hybrid deceptive strategy.

Keywords Multi-agent systems · Deception · Evolutionary computation

1 Introduction

Deception can play a role in any form of interaction, especially when the interaction involves two or more self-interested parties. Recent research has shown that in online interactions the temptation to deceive is particulary strong as users feel protected by anonymity (Castelfranchi 2001).

Open multi-agent systems are systems in which autonomous agents interact with one another to achieve some goal. This interaction can take the form of cooperation or competition. In open multi-agent systems in which agents are required to try and maximize their own utility, there can exist many opportunities and motivations to deceive (Ramchurn et al.

P. Ó Broin (✉) · C. O'Riordan
Department of Information Technology, National University of Ireland, Galway, Ireland
e-mail: pilib@zig.it.nuigalway.ie

C. O'Riordan
e-mail: colm.oriordan@nuigalway.ie

2004) Examples of this include: (i) agents within a system can represent different parties with different goals, some of these parties could be malicious, (ii) agents can join or leave a system at any time, potentially avoiding punishment or retribution for past wrong-doing, (iii) agents can have different abilities, roles or behavioural policies, (iv) agents typically have limited knowledge about their environment and must reason in the face of uncertainty; in order to do so agents sometimes rely on possibly deceptive information from other agents in the system.

Previous research on deception in multi-agent system in a simple blocks world called GOLEM was carried out by Castelfranchi et al. (1998), who showed that agents could successfully lie about their goals and abilities in order to obtain help from their credulous fellow agents.

Auctions have previously been used to examine the behaviour of rational trading agents (Wellman and Wurman 1999) and represent an interesting real-world model which has become increasingly popular in recent years due to the success of websites such as Ebay[1] and Amazon.[2] The English auction, also known as an ascending or open-cry auction is possibly the most commonly used auction protocol. Bidders announce their bids in increasing amounts until a single bidder remains. When this final bidder is left, the amount of their last bid is the price paid for the item. The dominant strategy for bidding in this type of auction is to bid slightly higher than the current bid until you reach your private valuation for the item. Given its widespread use, we adopt this as the auction protocol in our experiments.

In this work, we examine the potential shortcomings of two previously defined deceptive strategies (Ó Broin and O'Riordan 2006, 2007). We present a third deceptive strategy resulting from the combination of these two previous strategies and investigate its success relative to its predecessors. As in previous experiments, we measure not only the representation of the different strategy types in the test population over time, but also the average fitness of the society. This allows us to gauge the impact on societal fitness should deception spread among the population.

The next section briefly describes some common forms of deception possible in auctions. Section 3 deals with the experimental setup for both previous and current work including game design, Naïve strategy implementation and the evolutionary process used. In Section 4 we briefly give details on our previous work on deceptive strategies, outlining their encoding and providing a summary of results. Section 5 presents the motivation for the hybrid deceptive strategy and provides results of the experiments and their analysis. The final section includes a summary and general discussion and outlines possible future work.

2 Deception in auctions

Before we look at some different types of deception which can occur in auctions, we make the important distinction between private value and common value auctions. Private value auctions are auctions where bidders are not trying to acquire an item for resale or commercial use, but rather for personal consumption. Each bidder calculates their own private valuation i.e. how much the item is worth to them. These valuations can be similar or wide ranging. A particular coin, for example, is of little worth to a collector who has several similar coins, but may be worth a great deal to someone who needs it to complete their collection. Private value auctions are examined in this work.

1 http://www.ebay.com

2 http://www.amazon.com

In common value auctions, the item being auctioned has the same intrinsic value for all bidders even though bidders do not know what that value is. If, for example, a jar of coins is being auctioned, the jar has the same value for each bidder, but the amount they bid depends on their estimation of the value of the item.

Deception in real-world auctions is surprisingly common and can include techniques such as bidding rings, shill bids and sniping.

Bidding rings Bidding Rings represent one of the most common forms of bidder collusion in auctions. In a bidding ring, a number of bidders band together and agree not to outbid one another. This has the net effect of reducing the amount paid for the item. Once the designated member of the ring has won the item, the item is then re-auctioned among the ring members. The member of the ring who wins this second bid reimburses the bidder who won the first auction and then shares the remaining difference between the ring members.

Shill bids A shill bidder is a bidder who works for the auction house or seller who will bid either to get an auction started or if bidding slows in an auction. Recent research on the use of shilling on ebay can be found in (Barbaro and Bracht 2006) which discusses a process they call squeezing. This process makes use of the fact that bids can be cancelled to eliminate the risk to the seller in placing shill bids i.e. that they might win their own item. With squeezing, a seller makes use of a second user account on ebay to place shill bids designed to discover the private valuation of the highest bidder. Once they exceed the buyers valuation, they cancel the last bid and place a final shill bid just below the highest bidders reserve price. This ensures that the seller receives the maximum price achievable on the item.

Sniping Sniping refers to the process of late bidding in an auction. The question regarding the use of sniping was posed in (Roth and Ockenfels 2002) when the authors pointed out that the empirical evidence of late bidding in online auctions such as ebay which have a set auction end time, was at odds with the theory behind such auctions, which dictates that the timing of a bid is irrelevant for rational agents. The author in (Wilcox 2000) provides an explanation for such behaviour in the case of common value auctions. He notes that inexperienced bidders or those with little knowledge of an item can glean information about the value of an item from more experienced bidders or experts.

An explanation for sniping behaviour in private value auctions is provided in (Roth and Ockenfels 2005), in which the authors indicate that sniping can be a best response to an incremental bid strategy. When faced with an opponent bidder who places multiple bids, sniping not only avoids the possibility of a bidding war taking place, it also, if used effectively, negates the possibility of an incremental bidder being able to respond to a bid placed in the last few moments of an auction.

A further explanation for sniping is provided by Wang et al. (2004) who explore the use of sniping as a best response to dishonest sellers. In circumstances where a seller places shill bids, sniping minimizes the time which a dishonest seller has to cheat.

Antisocial bidding Antisocial bidding is another form of deception involving the placing of false bids. These bids are the equivalent of false signalling and are designed to deceive an opponent as to our level of interest in an item. Unlike shill bids placed by dishonest sellers as, for example, in (Barbaro and Bracht 2006), these bids are placed by bidding agents in an attempt to lower the profit of their fellow competitors. This kind of bidding is related to the notion of 'antisocial agents' as outlined in (Brandt 2000). In this work, the author comments on the fact that bidding agents are generally modelled as only being concerned for their own absolute profit, having no regard for the profit of their fellow bidders. He goes on to argue that there are some situations (such as in closed markets (Brandt and Weiss 2001)) in which agents are more interested in their relative profit i.e. their level of profit compared to that of their opponents. Brandt (2000) defines a formal payoff mechanism for agents incorporating a level of antisociality of the agent (how much value they place on reducing their opponent's profit compared to maximizing their own).

3 Experimental setup

This section examines the game model including the auction mechanism used, the Naïve (honest) strategy set design and the evolutionary process.

3.1 Game design

Agents compete in a repeated English auction with hard close (fixed end time). Each agent is assigned the same budget and a random subset of the complete set of auction items, this is referred to as the agent's *goal list*. Each item on this list is randomly assigned a portion of the agent's budget which corresponds to the agent's *private valuation* for that item. Each auction has ten timeslots in which bids can be placed on any item. At the end of an auction agents are assigned fitness scores based on their level of profit relative to the profit of their opponents.

A single game in the simulator consists of seven auctions, and the agents partake in 500 games per generation. Different goal lists are assigned to the agents in each game, ensuring that strategies must be able to compete effectively across a wide variety of auction scenarios.

3.2 Strategy design

We begin by defining four Naïve bidding strategies. The behaviour of these Naïve strategies is encoded by three genes *random, linear* and *dynamic*. These are termed *naïve genes*.

Fixed For a strategy to be fixed, all three of its naïve genes must have a zero value. Using a fixed strategy, an agent's bid for an item is calculated as a percentage of their private valuation for that item determined by the amount of time elapsed in the auction.

Random The random gene in an agent's genotype indicates the probability that the agent's bid will be random (a value of 0.7 indicates a 70% chance of a random bid). This means that 70% of the time, an agent's bid is calculated as a random value between the current bid on the item and the agent's private valuation for that item. Should this occur, the strategy is termed purely random and the values for the rate and dynamic genes are equivalent to 'don't cares'. For the remaining 30% of the time, the random gene is interpreted as having a value of zero and the agent's bid is determined solely by the values of the remaining two naïve genes. This gives rise to a subset of random strategies incorporating random linear, random dynamic and random fixed.

Linear With a linear strategy an agent's bid is calculated as:

$$bid = curr(j) + (pv(i, j) - curr(j)) * rate \tag{1}$$

where $pv(i, j)$ indicates the private valuation of agent i for item j, $curr(j)$ is the current bid on item j and rate is the value of the linear gene in agent i's genotype. With a linear strategy, the rate is increased by 0.1 after each bid, leading to more aggressive bidding as the auction progresses.

Dynamic The dynamic bidding strategy is quite similar to the linear strategy; the main difference lies in the way in which the rate is updated. In the linear strategy, the rate is increased by 0.1 after each bid; with the dynamic strategy this value can vary. A dynamic gene value lower than 0.1 means that an agent will reach their private valuation for an item more slowly than they would with a linear strategy while conversely a higher value will mean that they reach it more quickly. This gives rise to varying levels of aggressiveness in bidding.

These naïve strategies are evolved to produce a fit test population into which our deceptive agents can be introduced.

3.3 Evolutionary process

Within our model, once the established number of games for a particular generation of agents has been played, the sum of profit for all of the agents is calculated. Any given agent's fitness for a particular generation is then given by the following equation:

$$fit(A_i) = \frac{prof(A_i)}{\sum_{j=1}^{N} prof(A_j)} \tag{2}$$

Once agents have been assigned a fitness value, they are sorted based on this value. Next, an agent is assigned a rank based on their position in the sorted fitness array. These ranks, which dictate an agent's probability of being selected for crossover are given by (Pos in Fitness Array)^{SP}, where SP is the Selective Pressure for the run. This form of linear ranking helps emphasize the difference between agents when fitness variance is low, and ensures that convergence does not happen too quickly when fitness variance is high. The selective pressure plays a vital role in the selection process, and even slight variation in its value can cause a huge bias towards fitter strategies in the population. It is worth noting that in our model, if an agent has a zero fitness score, then its probability for selection is zero.

Once the ranks have been assigned, a process of roulette wheel selection is begun. The higher the rank an agent has, the larger the portion of the roulette wheel associated with that agent is. Once parents have been selected, they are copied to the mating pool where the process of crossover can begin.

There are two different types of crossover in our model: weighted and standard. The weighted crossover represents a gradual progression towards fitter strategies, and can be viewed as local search in the fitness landscape. In this method of crossover, each parent is assigned a crossover-weight based on its fitness. The impact which each parent has on the genetic makeup of the offspring is determined by this weight; the offspring is more similar in genotype to the fitter of its two parents (Fig. 1).

Although allowing for faster convergence when dealing with real-valued geneotypes, this type of crossover will only produce a genotype that lies somewhere between that of the two parents; it is therefore necessary to include a measure of standard crossover which presents the possibility of combining favourable gene values from both parents to produce an individual which is fitter than either of its parents. In the current version of the auction simulator, weighted crossover occurs with a probability of 70% while standard crossover occurs the remaining 30% of the time. It was found that this combination allows convergence within

	gene 1	gene 2	fit.
Parent A	.95	.34	.05
Parent B	.67	.12	.03

	gene 1	gene 2	fit.
Child	.845	.2575	?

Fig. 1 Weighted crossover

a reasonable amount of time while avoiding scenarios in which strategies quickly become caught in local peaks.

As our model deals with strategies represented by real numbers, there was also a need to alter the standard model of mutation used in genetic algorithms which deal with binary representations. Rather than the flipping of a 1 to a 0 or vice versa, our model probabilistically increments or decrements a real number by a value within a given range. The current range is [0, 0.3], and the probability of mutation occurring stands at 3%. Unlike mutation in a binary representation, where flipping a 0 to a 1 can give a vastly different strategy (e.g. 0 0 0 1 to 1 0 0 1), our mutation does not have such a huge impact on the variety of strategies in the population. Again, this is accounted for by the 30% probability of standard crossover occurring.

Our evolutionary model also includes a small measure of elitism whereby the top 5% of strategies are directly copied from generation m to generation $m+1$ without being subjected to the normal processes of crossover and mutation.

The next section introduces our two previous deceptive strategies and their encoding as well as a summary of some past results.

4 Previous work

4.1 Sniping and antisocial bidding

In previous work (Ó Broin and O'Riordan 2006, 2007) we have shown how auctions can be used as a simple mechanism for investigating notions of deception in a multi-agent system. We have examined two forms of deceptive behaviour which can evolve in a repeated English auction setting—sniping and antisocial bidding.

Sniper The goal of our sniping agent was to hide information about its true valuation for an item, while simultaneously allowing it to win items at a price below that which it would normally be forced to pay had a 'Naïve' (honest) strategy been employed. Sniping also allowed agents to win items for which they do not possess the highest valuation. Under normal circumstances, the bidder with the highest private valuation would be expected to win an item since they will pay the most to win that item. With sniping, the agent with the highest valuation for an item will hold the current high bid on that item until just before the auction's end. A sniper will then place a bid in the last few moments of an auction. If the auction were to continue, the agent with the highest valuation would then simply increases their bid, but because the auction ends, they do not receive the chance to increase their bid and thus the sniper wins the item even though it has a lower valuation than one of its opponents.

Antisocial bidder The behaviour of our antisocial bidding agent (hereafter referred to as *Dec2*) was based on the notion of relative fitness (concentrating not only on maximizing one's own fitness, but also on minimizing that of other agents), *Dec2* bidders place 'false bids' (bids for items which they do not wish to acquire) which serve to ensure that, over time, an opponent will be forced to pay an amount increasingly close to their private valuation in order to win that item. This has the net effect of reducing the payoff for the agent who wins that item[3] and thus reducing their chance of survival.

4.2 Encoding and previous results

The behaviour of a strategy defined as being deceptive is based on three *deceptive genes*. The first deceptive gene (*dec*), represents the type of deception used by the agent. In experiments with a single type of deceptive agent, a value for this gene of less than 0.5 indicates a Naïve strategy, while anything above this value indicates a deceptive one. When more than one type of deceptive strategy is involved in an experiment, the dec gene range ([0..1]) is split three ways, allowing us to encode each of the three types of deceptive strategy—Sniper, *Dec2* and Hybrid.

The remaining two deceptive genes code for the behaviour of the Sniper (or Hybrid strategy which incorporates sniping). These genes (*bid-time* and *bid-amount*), control how late in the auction bids are placed (i.e. a bid-time of 0.7 indicates bidding in the seventh timeslot), as well as the value of those bids. In order to accurately portray real world sniping behaviour, we would expect bid-time to evolve to a high value while the bid-amount gene evolves to a low one.

The *Dec2* strategy ignores the bid-time and bid-amount genes and instead uses its naïve genes for bidding on items on its goal list. The strategy makes use of the information gained in a repeated English auction about the private valuations of its opponents to increase its own profit while minimizing the profit of its fellow bidders. To do this, it incorporates two additional data structures (*high-bid* and *low-bid*) which it uses to maintain a limited bid history for auction items.

The first of these structures is used to store the highest bids placed so far on each item which the agent does not win. Once an agent has a high-bid value greater than zero stored for a particular item, it will bid that amount (less a small reduction) for the item in the next auction. This means that the agent which won the item in the previous auction must increase their bid this time around. This process continues with the values stored for each item in the high-bid structure increasing until they reach the private valuation of the highest bidder for that item. This means that the potential profit for other bidders on this item has effectively been reduced to zero.

In cases where an agent wins an item, the low-bid data store is used to record the lowest bid with which the agent has thus far won that item. In these cases the high-bid data structure is used to store the highest bid at which an opponent has dropped out of the bidding for that item (i.e. the second highest bid for that item). In subsequent auctions, the agent will bid somewhere between these two values thereby increasing its profit. As the auctions progress, these two values will be adjusted based on new bid information, eventually settling just above the second highest bid for the item. The agent can then bid this amount knowing that it will win the item for the lowest possible price.

Previous results showed that antisocial bidders effectively reduce the payoffs of competitors to zero on many items (Ó Broin and O'Riordan 2006). The second part of the *Dec2*'s bidding behaviour was designed to allow it to maximize its own profit. Having won an item

[3] Payoffs are calculated as the difference between an agent's private valuation for an item and the price actually paid to win that item.

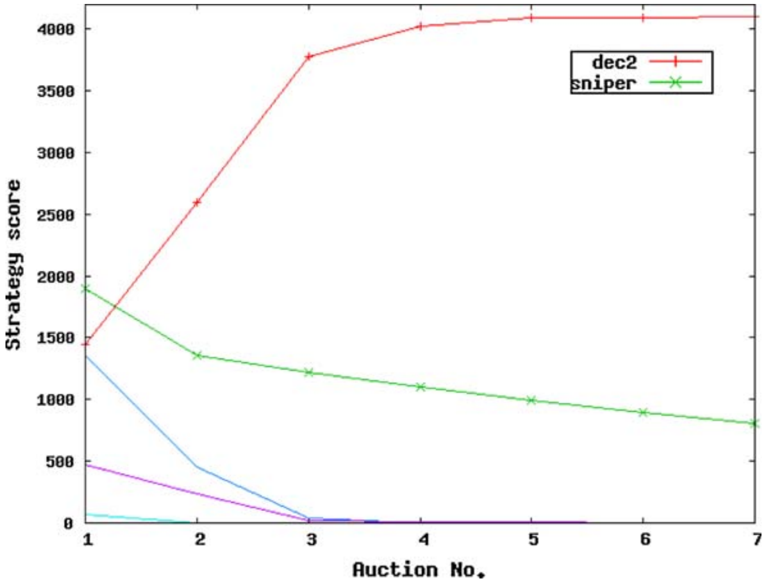

Fig. 2 Effect of *Dec2* on Naïve and Sniper profit

in an auction the *Dec2* agent gradually reduces the amount which it will pay for that item in future auctions thus increasing its profit. As the auctions continue, the *Dec2* agent's bid for that item will eventually settle at a value just above the highest bid of the agent with the next highest private valuation for that item. The *Dec2* bidder thus achieves the maximum available profit on that item.

Both the sniping and *Dec2* strategies were shown to achieve higher payoffs than the test population of Naïve strategies into which they were introduced. These high levels of payoff ensured that the deceptive strategies were selected for reproduction and quickly spread among the Naïve agents. In co-evolutionary settings, snipers were seen to be somewhat robust to the profit-reducing effect of the *Dec2* strategy and populations usually converged to the former within 15 generations. As we see in Fig. 2, the *Dec2* strategy succeeds not only in increasing its own profit but also in dramatically reducing the profit of the three Naïve agents in the population. Although it also reduces the profit of the sniper it does so on a much more gradual basis since the sniper reveals less information about its private valuations for the items on auction. This result coincides with the findings in (Barbaro and Bracht 2006), which indicate that sniping is a possible best response to the real-world scenario of 'squeezing'.

5 Hybrid strategy motivation, results and analysis

Although both types of strategy successfully exploit Naïve populations, they each implement only one form of deception. Snipers, while winning items for which they do not possess the highest private valuation, do nothing to reduce the profits of their opponents. This is a missed opportunity for increasing their relative fitness.

Dec2 bidders, on the other hand, focus mainly on decreasing their opponents' profit, they do not try to use late bidding in order to win items they wish to acquire at lower prices. A successful hybrid strategy would need to incorporate both types of deception in order to fully

Table 1 Hybrid position in Naïve fitness rankings

Rand	Linear	Dyn	Bt	Ba	Dec	Fitness
0.47	0.95	0.05	0.55	0.42	0	0.0183
0.25	0.80	0.51	0.27	0.29	0	0.0310
.	.	.	.	.	.	.
.	.	.	.	.	.	.
0.39	0.26	0.85	0.03	0.76	0	0.0742
0.07	0.03	0.29	0.82	0.32	0	0.1749
0.14	0.58	0.44	**1**	**0.1**	**0.75**	**0.4280**

exploit the weaknesses of its opponents. A hybrid strategy in our simulator (indicated by its deceptive gene value) therefore utilises sniping (bid time, bid amount genes) for items on its goal list while simultaneously placing false bids (using information stored in bid array structures) on items which it does not wish to acquire.

The Hybrid strategy is tested against our three previous strategies (Naïve, Sniper and *Dec2*). We begin by testing the performance of the Hybrid strategy against Naïve bidders in both non-evolutionary and evolutionary settings. We then show the results following a single generation of introducing one Hybrid strategy into populations of Snipers and *Dec2*s. Finally, we examine the performance of the Hybrid strategy in a population consisting of an equal mix of all three deceptive strategies which is allowed to evolve over 500 generations.

5.1 Hybrid vs. Naïve

Non-evolutionary setting We begin by testing our Hybrid strategy against Naïve bidders in a non-evolutionary environment to ascertain its position in the fitness rankings after one generation. The test population consists of nine randomly initialized Naïve agents and one Hybrid agent playing 500 games with a goal list of five items each from a complete auction set of ten items. The Hybrid agent has a deceptive gene value of 0.75 while the Naïve strategies' deceptive genes are set to zero. The Hybrid strategy has a bid-time of 1 (indicating that it will bid in the last possible timeslot), and a bid-amount of 0.1.

As we can see in Table 1, the Hybrid strategy (bold) receives a fitness score 0.42 while the next fittest strategy receives a score of 0.17. This fitness is higher than a Sniper (0.21) and comparable to a *Dec2* agent (0.43) in a similar scenario (Ó Broin and O'Riordan 2006).

In a population of fit (pre-evolved) Naïve strategies, the Hybrid's ability to reduce the profit of its opponents is limited due to the Naïve strategies' use of small bid increments which cause them to reach their true valuations more slowly. The Hybrid strategy's use of sniping, however, ensures that it still consistently achieves the highest fitness score.

Evolutionary setting Next we examine the representation of the different strategies in the population as well as the societal fitness over 500 generations for the test population used above.

Figure 3 shows the number of Naïve and Hybrid strategies in the population at each generation. The Hybrid strategy quickly spreads and dominates the population for the remainder of the run despite a few Naïve strategies reappearing through mutation.

The effect of this spread of the Hybrid strategy on societal fitness is seen in Fig. 4. The sharp rise in the first ten generations is a result of more and more of the agents in the

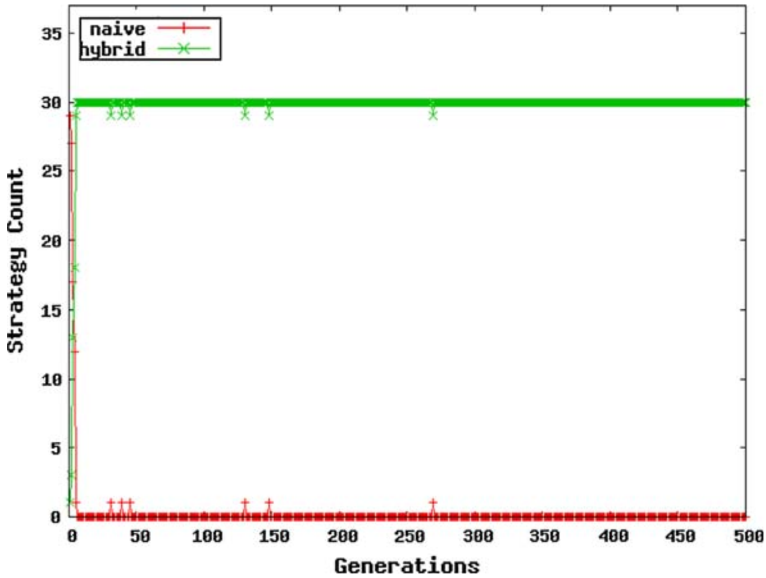

Fig. 3 Strategy count with Naïve and hybrid strategies

population adopting a Hybrid approach to bidding. As previously seen in the case of snipers, late bidding means a decrease in the amount paid to win an item and thus a corresponding increase in profit and overall societal fitness. The remainder of the run sees the societal fitness fluctuate as the bid-amount gene undergoes a process of genetic drift as outlined in (Ó Broin and O'Riordan 2006). This can be explained as follows.

The placement of snipers in the bidding order at any given timeslot in an auction is a function of their bid-time. Snipers with a bid-time of 0.95 will bid earlier than those with a bid-time of 0.98 and so on. Snipers with the same bid-time are randomly ordered within their position in the bidding order, so three snipers bidding with a bid-time of 0.98 will bid in a random order after those with the 0.95 bid-time. When all sniper bid-times quickly converge on the optimal 1.0 (indicating bidding at the last possible moment in the auction), the success of an individual sniper in the population is largely independent of their bid-amount. This causes the bid-amount gene to drift from higher to lower values and vice versa causing fluctuations in the societal fitness.

5.2 Hybrid vs. Sniper

Table 2 shows the fitness rankings after a single generation (500 games) when one Hybrid strategy is introduced into a population of nine Snipers. In this experiment the deceptive gene was used to encode the three different deceptive strategies (Snipers > 0, *Dec2* > 0.33, Hybrid > 0.66). As such, the strategies were assigned a deceptive gene value roughly halfway through the range of values encoding their individual behaviour (i.e. 0.17 for Snipers and 0.83 for Hybrids). Both the Snipers and Hybrids had a bid-time of 1 and a bid-amount of 0.1.

The Hybrid strategy achieves a slightly higher fitness score (0.109) than the fittest Sniper (0.105). The sniping of the Hybrid strategy puts it on a par with the Snipers in terms of

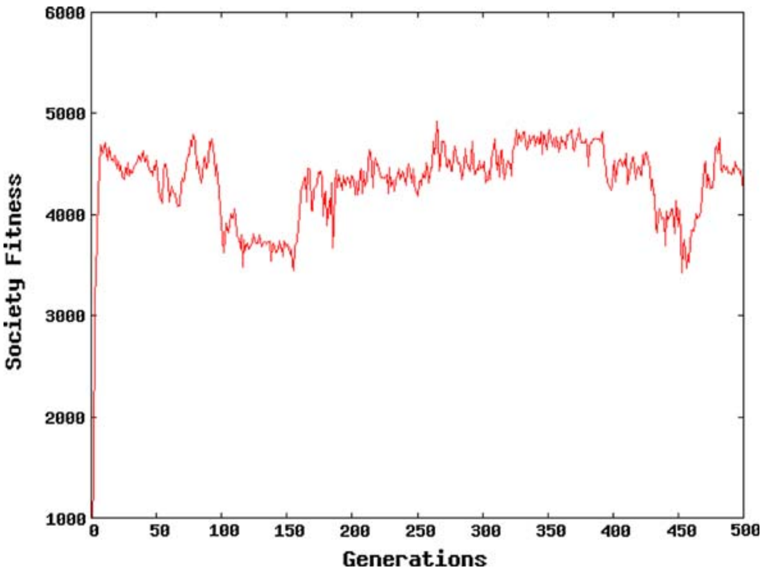

Fig. 4 Societal fitness with Naïve and hybrid strategies

Table 2 Hybrid position in Sniper fitness rankings

Rand	Linear	Dyn	Bt	Ba	Dec	Fitness
0.48	0.69	0.11	1	0.1	0.17	0.0920
0.24	0.73	0.79	1	0.1	0.17	0.0967
.	.	.	.	.	.	.
.	.	.	.	.	.	.
0.12	0.64	0.58	1	0.1	0.17	0.1024
0.75	0.62	0.17	1	0.1	0.17	0.1057
0.51	0.35	0.66	**1**	**0.1**	**0.83**	**0.1094**

winning items for which it has a private valuation >0, while the *Dec2* behaviour of placing false bids reduces its opponents' profit. The net result is that the Hybrid strategy wins items at a rate similar to Snipers, but reduces the Snipers' profit while its own remains unaffected thus achieving a higher relative fitness.

5.3 Hybrid vs. *Dec2*

In these experiments, *Dec2* strategies were assigned a deceptive gene value of 0.5, while Hybrid strategies were again encoded by a value of 0.83. The Hybrid strategy had a bid-time of 1 and a bid-amount of 0.1. The *Dec2* strategy had its Naïve genes set to the final values which were observed in Naïve strategies which were evolved for 500 generations.

As we would expect the *Dec2* strategies have the usual effect on one another (and on the Hybrid strategy) of decreasing an opponent's profit. The Hybrid strategy, however, achieves a higher fitness score due to its use of sniping. As previously shown, the *Dec2* strategy is

Table 3 Hybrid position in *Dec2* fitness rankings

Rand	Linear	Dyn	Bt	Ba	Dec	Fitness
0	0.02	0.02	0.65	0.05	0.5	0.0871
0	0.02	0.02	0.24	0.60	0.5	0.0888
.	.	.	.	.	.	.
.	.	.	.	.	.	.
0	0.02	0.02	0.90	0.17	0.5	0.0979
0	0.02	0.02	0.09	0.30	0.5	0.0991
0.52	0.28	0.67	**1**	**0.1**	**0.83**	**0.1536**

Table 4 All three deceptive strategies—fitness rankings

Rand	Linear	Dyn	Bt	Ba	Dec	Fitness
0	0.02	0.02	0.20	0.78	0.5	0.0124
0	0.02	0.02	0.56	0.22	0.5	0.0135
0	0.02	0.02	0.12	0.56	0.5	0.0145
.	.	.	.	.	.	.
.	.	.	.	.	.	.
0.56	0.95	0.46	1	0.1	0.83	0.0430
0.12	0.92	0.32	1	0.1	0.17	0.0432
0.68	0.24	0.31	1	0.1	0.17	0.0439
0.42	0.78	0.22	1	0.1	0.17	0.0440
0.89	0.85	0.19	1	0.1	0.83	0.0445
0.21	0.01	0.37	1	0.1	0.17	0.0463
0.12	0.72	0.45	1	0.1	0.83	0.0468

not as effective in reducing the profit of Snipers as it is in the case of Naïve bidders. This means that the sniping behaviour of the Hybrid strategy gives it an advantage over the *Dec2* strategies in this scenario.

5.4 All three deceptive strategies

This section provides the results of an experiment involving a population of 30 agents consisting of 10 *Dec2*, 10 Sniper and 10 Hybrid strategies.

In the fitness rankings after one generation, the *Dec2* strategy type (deceptive gene value of 0.5) performs the worst, with the ten seeded strategies achieving the ten lowest scores. The Sniper (deceptive gene value of 0.17) and Hybrid (deceptive gene value of 0.83) strategies are randomly ordered in the top 20 positions. This may seem odd considering that the Hybrid strategy achieved the highest score when seeded in a population of snipers, but can be explained by the presence of further Hybrid (as well as *Dec2*) strategies.

When our Hybrid strategy beat each of its sniping competitors in the non-evolutionary setting, it was the only strategy which was placing false bids to decrease it opponents' profit. Since there is more than one strategy incorporating this kind of bidding in this new population,

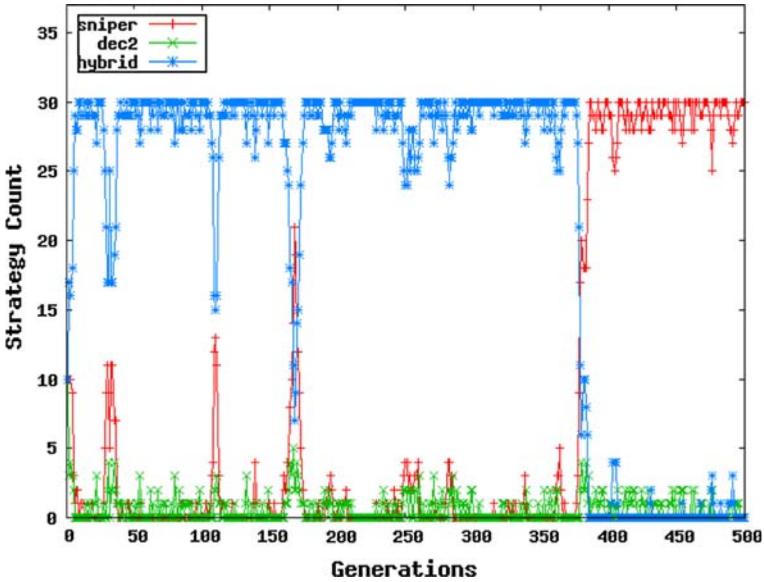

Fig. 5 Strategy count with all three deceptive strategies

the Hybrid strategies also have their profit reduced. This in effect negates the benefit of the *Dec2* component of the Hybrid strategy, leaving its success dependent solely on its sniping genes. In essence, this means that a Hybrid strategy in a population with fellow Hybrids, or indeed with *Dec2*s will perform as though it were simply a Sniper.

Since snipers with the same bid-time place bids in a random order,[4] their success in winning items and thus their fitness is also random. This accounts for the mix of Sniper and Hybrid strategies in the top 20 positions in the fitness rankings. The effect that this has on the representation of the various strategies in the population over time is shown in Fig. 5.

Since the *Dec2*s consistently achieve the lowest fitness scores, they die off within the first ten generations and do not gain a foothold for the remainder of the run. As the fitness values of the Sniper and Hybrid agents are essentially the same, the population does not fully converge to either of these strategies, but instead fluctuates between the two. In the example shown the Hybrid strategy places highest in the fitness rankings after the first generation and so dominates initially. If mutation introduces enough snipers, or if several snipers are highly placed in the fitness rankings, then a 'critical mass' is reached in which the evolutionary pressure switches from Hybrids to Snipers (as happens at approximately generation 375). This cycle will continue as long as the population is allowed to evolve and ultimately neither strategy will remain indefinitely stable.

The societal fitness for this population initially increases as the *Dec2* strategies change to either the Sniper or Hybrid strategies within the first 10–15 generations. After that point the average fitness fluctuates similarly to that of a population consisting entirely of snipers, although the inclusion of some *Dec2*s sees a slight reduction in the average fitness due to their decreasing of opponents' profits (Fig. 6).

[4] The equivalent of 30 people placing a bid in an online auction at the exact same time.

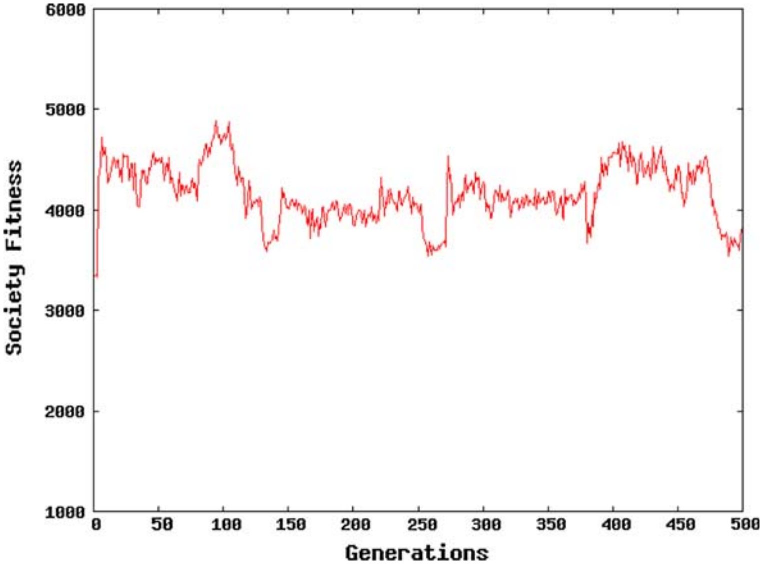

Fig. 6 Societal fitness with all three deceptive strategies

6 Summary and conclusions

Understanding issues relating to trust and deception is vital not only in the field of multi-agent systems, but also in areas such as sociology, political science and business, or indeed any other field which involves predicting or reasoning about the behaviour of self-interested parties.

The examining of these concepts can lead not only to better understanding of their function in a social context, but also, in technological terms, to the design of more reliable systems which would be robust to misuse by malicious behaviour. Any multi-agent system which is designed to constrain deceptive behaviour is inherently more likely to elicit a higher level of trust and confidence in its users.

This paper has examined the limitations of previously discussed deceptive strategies for the repeated English auction. We have proposed a Hybrid deceptive strategy incorporating successful elements from both of these strategies and tested its performance against both Naïve strategies and previous deceptive strategies in an evolutionary setting. The Hybrid strategy was found to have advantages over the Sniper and *Dec2* strategies individually over a single generation, but suffered from the same problem as *Dec2* in an evolutionary setting i.e. when faced with an opponent using the same profit-reducing technique, its performance was diminished. This being said, the Hybrid strategy proved more successful than the *Dec2* strategy and equally as successful as the Sniper in evolutionary settings, with populations converging to the Hybrid strategy in 50% of experiments and to the Sniper strategy in the remaining 50%. The experiment can also be seen as a confirmation of the robust, effectiveness of the Sniper strategy, which can be seen as having the advantages of less complex behaviour and fewer resource requirements.

Future work will focus on examining the effect of varying the evolutionary operators and their rates as well examining the types of deception which can occur in alternative auction protocols. We will also explore the possibility of allowing *Dec2*/Hybrid strategies to recognize one another and thus implement bidder collusion in an attempt to address their shared weakness.

Acknowledgements The authors wish to thank the anonymous reviewers for their helpful comments.

References

Barbaro S, Bracht B (2006) Shilling, squeezing, sniping: explaining late bidding in online second-price auctions. Working paper, University of Mainz (this version: January 31, 2006)

Brandt F (2000) Antisocial bidding in repeated Vickrey auctions. Tech. Rep. FKI-241-00. Department of Computer Science, Technical University of Munich, ISSN 0941-6358

Brandt F, Weiss G (2001) Antisocial agents and vickrey auctions. In: Meyer JJC, Tambe M (eds) Intelligent agents VIII. Lecture notes in artificial intelligence (LNAI), vol 2333, revised papers from the 8th International Workshop on Agent Theories, Architectures and Languages (ATAL). Springer-Verlag, pp 335–347

Castelfranchi C (2001) The role of trust and deception in virtual societies. In: Proceedings of the 34th Hawaii international conference on system sciences

Castelfranchi C, Falcone R, de Rosis F (1998) Deceiving in golem: how to strategically pilfer help. In: Deception, fraud and trust in multiagent systems. Kluwer Publishing

Ó Broin P, O'Riordan C (2006) Deception in multi-agent auctions: An evolutionary approach. In: 17th Irish conference on artificial intelligence and cognitive science (AICS2006)

Ó Broin P, O'Riordan C (2007) Evolving a hybrid deceptive strategy for the repeated english auction. In: Proceedings of the 18th Irish conference on artificial intelligence and cognitive science (AICS2007)

Ramchurn SD, Huynh TD, Jennings NR (2004) Trust in multi-agent systems. Knowl Eng Rev 19(1):1–25

Roth AE, Ockenfels A (2002) Last-minute bidding and the rules for ending second-price auctions: evidence from ebay and amazon auctions on the internet. Am Econ Rev 92:1093–1103

Roth AE, Ockenfels A (2005) Late bidding in second price internet auctions: theory and evidence concerning different rules for ending an auction. Games Econ Behav 50

Wang W, Hidvegi Z, Whinston AB (2004) Shill-proof fee (spf) schedule: the sunscreen against seller self-collusion in online english auctions. Goizueta Paper Series, Emory University

Wellman MP, Wurman PR (1999) A trading agent competition for the research community. In: Proceedings of the international joint conference on artificial intelligence. Workshop on agent-mediated electronic trading. AAAI Press

Wilcox RT (2000) Experts and amateurs: the role of experience in internet auctions. Market Lett 11:363–374