Distributed Constraint Satisfaction Workshop,
First International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-2002), Bologna, Italy, July 15-19, 2002 (http://lia.deis.unibo.it/aamas2002/)

Simple Constrained Bargaining Game

Edward Tsang and Tim Gosling
Department of Computer Science, University of Essex, Wivenhoe Park, Colchester CO4 3SQ, UK
Phone: +44 01206 87{2771,2774,2787}
E-mail: {edward, tdbgos}@essex.ac.uk
URL: http://cswww.essex.ac.uk/CSP/

2 May 2002

Abstract. Motivated by the need to solve complex bargaining situations in real life applications, this paper defines a Simple Constrained Bargaining Game, CBG(S), which extends the games that have been tackled in mathematical game theory so far. In the game, the seller is constrained by a cost and the number of days to sell; the buyer is constrained by a utility and the number of days to buy. Neither side knows the opponent's constraints. This is a simple problem from the real life sense, but hard for mathematical analysis. Strategies for tackling this problem have been examined. Some of these strategies, together with a platform for conducting computational experiments in this game, have been implemented. Preliminary results are reported. Components in these strategies are identified to prepare for an evolutionary approach that tackles the Simple Constrained Bargaining Game.

Keywords: bargaining, constraints, game, evolutionary computation

1 Introduction

This work is motivated by realistic applications. We focus on a fundamental bargaining problem, which extends the problems attacked by game theorists so far.

With the growth of activities on the Internet, more machines have to communicate with each other [11]. Computer bargaining is quite different from human bargaining, where psychology, emotion, (in the case of audio communication) tone, (in the case of face to face bargaining), eye contact, appearance and facial expressions could all play important parts. There is a need to develop automatic (algorithmic) bargaining agents.

Game theoreticians have long been studying bargaining strategies [1, 4, 9]. Optimal strategies have been found for many problems. The game theoretic approach is attractive because it is neat, and results often have provable properties.

However, realistic applications are far more complex than those studied by game theoreticians so far [7, 8]. Jennings et al [8] quite rightly pointed out the limitation of game theoretic models, due to complexity. By using computer programs and simulations, complex environments and ideas in game theory can be evaluated and verified. This is the approach taken in this research.

2 The Simple Constrained Bargaining Game

In a supply chain, suppliers, middlemen and consumers negotiate with each other. They compete with each other for profit. They also cooperate with each other in order to exploit opportunities. We focus on the situation where one buyer negotiates with one seller only, in the hope of making a transaction.

The simple constrained bargaining game, CBG(S):

The game involves a buyer and a seller.

- (i) The seller is constrained by a *cost* and the number of days within which it has to sell (*DTS*);
- (ii) The buyer is constrained by a *utility* and the number of days within which it has to buy (*DTB*);
- (iii) The seller does not have any information about the buyer's utility and DTB;
- (iv) The buyer does not have any information about the seller's cost and DTS;
- (v) The players make alternative bids, with the seller to bid first;
- (vi) Each player bids exactly once per day;
- (vii) When both players bid for the same price, a sale is agreed;
- (viii) If a sale cannot be agreed before a player runs out of time, the negotiation terminates; no penalty is paid by either player;
- (ix) This is a one-off game: neither player has information about the others' past behaviour and performance.

For example, the seller may have a cost of 100 and it may have to sell within 12 days. The commodity may have a utility of 500 to the buyer, who has to buy within 16 days. The seller may start the bid at, say, 215. The buyer may counter offer 53, which the seller counter offer 192, and the bargain continues until, say, the seller asks for 120 on or before day 12 which the buyer accepts.

Success of an agent in one game depends to some extent on luck. However, if enough agents play many games against each other, some strategies may prove to be better than others. Success of a strategy can be measured by comparing its profit against those by other strategies. Therefore, performance of a strategy depends on what other strategies are involved in the game. The question is whether optimal strategies exist, and if so, what they look like.

This is an interesting problem because it extends what has been done in game theory to something that is still manageable. Given enough effort, it may even be possible to find evolutionary stable solutions. It is also interesting because it forms the basis of more realistic situations which involve more constraints than just cost, utility, sell by dates and buy by dates.

3 Strategies for CBG(S)

When the utility of the buyer is higher than the cost of the seller, there is opportunity for both players to make a profit. This is the incentive for the agents to strike a deal in CBG(S). For convenience, we call the range between the cost and the utility the "profitability region".

Within the profitability region, agents would try to maximise their profit. This is the incentive for the buyer (seller) not to make a deal even when the asking (bid) price is below the utility (cost). Any strategy should balance between not missing profitable opportunities (when utility is above cost) and not conceding too much profit to the opponent.

Following are three first attempt, non-trivial, buyers and sellers that have performed reasonably well against naïve opponents:

The Keen-Seller-2

Summary: This seller is keen to make deals, but when time is available, it attempts to get a better deal by delaying commitment by one round.

(A) First offer: Cost plus 1 times a multiplier, which is equal to DTS/2 or 2, whichever greater.

(B) When the buyer's offer is above its cost:
The bid will be accepted if and only if:
   (a) It has no more than M days to spare, where M is a parameter; and
   (b) The previous offer was already above the cost.

When the bid is not accepted, the Keen-Seller will offer the difference between the cost and the previous offer divided by the number of days left.

(C) When the buyer's offer is below its cost:
Offer half way between the cost and the previous offer.

The Smart-Seller-4

Summary: A *Target* is worked out, principally based on an estimation of the pattern of the buyer's previous bids. Up to three bids are used to project the buyer's next bid.

(A) First offer: Bid cost times a multiplier, which is equal to DTS/2 or M (which is a parameter), whichever greater.

(B) When the buyer's offer is above its cost:
The bid will be accepted if and only if:
   (a) It has no more than D days to spare, where D is a parameter; and
   (b) The bid price is within R% of the Target, where R is a parameter.

When the bid is not accepted, the Keen-Seller will offer the difference between the cost and the previous offer divided by the number of days left.

(C) When the buyer's offer is below its cost:
Offer the AP – (AP – Target)/DL, where AP is the seller's previous asking price and DL is the number of days left to sell.

The Progressive-Buyer-2

Summary: The idea is to divide the utility value by the DTB, increasing the bid gradually. Therefore, if the utility is 720, and there are 6 days to buy, the general rule, which could be overridden, is to bid 120, 144, 180, 240, 360 and 720.

(A) First offer: if the seller's asking price is below utility, bid the asking price divided by DTB; else proceed with the general rule

(B) When the seller's offer is below utility:
The bid will be accepted if and only if:
   (a) It has no more than D days to spare, where D is a parameter; and
   (b) The ratio between the last two asking prices by the seller is within a certain ratio R, where R is a parameter.

When the offer is not accepted, repeat the previous bidding price

(C) When the seller's offer is above its utility:
Bid the Utility divided by the number of days left.

More strategies have been implemented, including strategies for control experiments:

- A random seller, which starts with a high offer, then offers a random value between the cost and the last bid. A buyer of similar behaviour was also implemented.
- Keen-Seller-1, which behaves exactly like Keen-Seller-2 above, except that it accepts any bid that is above its cost.
- A number of smart buyers that use similar but different rules from Smart-Buyer-2.
- A progressive seller that mirrors the behaviour of the Progressive-Buyer-2.
- A smart buyer that mirrors the behaviour of the Smart-Seller-2.
- Over a dozen buyers and sellers which are yet to be fully evaluated.

4 Experimentation

A Mediator (Version 1.9) calls one seller and one buyer in turn and feed them with their opponents' offers. It then summarises the results of a given number of games between these two opponents. A Game Platform (Version 2.9) is used for conducting large-scale experimentation among a set of sellers and buyers. Results are summarised in a table form. Tables 1 and 2 show the result of two typical runs.

It is worth noting that it is very difficult to fully evaluate the performance of a strategy. In CBG(S), as it is the case in many other games, the performance of a strategy depends on:

- Who it is playing against; and
- The parameters used by the experimenter (which is different from the parameters used by the strategy itself) to run the experiments.

For example, Table 1 shows that Smart-Seller-4 had the best performance among the sellers, with an average score of 84.8 per game. Keen-Seller-2 scored on average 79.9 per game. However, had the two "progressive buyers" not been present in the game, Keen-Seller-2 would have out-performed Smart-Seller-4 (70.0 versus 68.8). This is because the "progressive buyers" scored more heavily on the Keen-Seller-2 than they do on Smart-Seller-2.

The relative strength of the strategies also changes when the size of the profitable region (utility minus cost) changes. Table 2 records the experiments with a narrower profitable region. In Table 2, Keen-Seller-2 scored 51.4 per game on average, out-performing Smart-Seller-4. Careful scrutiny reveals that although Smart-Seller-4 obtained on average 48% of the profits in the profitable region (Table 3(b), last column), it only managed to realise 70% of the potential profits together with its opponents (Table 4(b)). Keen-Seller-2 managed to realise 87% of the potential profits. In fact, Keen-Seller-2 did not lose out to its opponents: it took 51% of the profit against its opponents (mainly because there is less profit for the progressive buyers to exploit).

Progressive-Buyer-2 seems to perform well against all the sellers tested so far. Having said that, more sophisticated buyers and sellers have been developed. They perform better than the Keen-Seller-2, Smart-Seller-4 and Progressive-Buyer-2 under the parameter settings tested so far. However, their properties are yet to be studied, and therefore they will not be presented here.

5 Strategy Templates: Components of Strategies

We have presented the strategies above under template, which define the decision components of a strategy. More complex templates are being studied, which includes more refined strategies close to the final day. We shall refer to such templates as *strategy templates*, which describes what the agent does under each situation.

There are a limited number of obvious rules under each entry of the strategy template. Having written down the decision components and possible rules to apply, we are in a position to explore the possibility of generating templates automatically. This involves (a) instantiating each slot in the template with a rule; and (b) determining parameter values.

We are attempting to use an evolutionary approach to generate strategy templates [6, 5, 9]. The Game Platform will be used to provide feedback to the performance of a strategy template, which will hopefully guide the search for fitter templates.

We have seen above that the performance of a strategy depends on the behaviour of other participants. One of the attractiveness of adopting an evolution approach is that, when past experience and market information are available, agents can switch strategies in adverse situations.

6 Conclusions and Future Work

In this paper, we have focussed on a simple constrained bargaining game, CBG(S). The problem is analysed. A heuristic approach has been taken and a number of hand-coded agents have been implemented. A Game Platform has been built for evaluating agents. Experimental results so far suggest certain useful strategies. New, more effective strategies are in the pipeline. Our experience so far enables us to write down the decision components for participating in CBG(S). These components will be used to develop evolving agents in the future.

This research is part of a larger project with business motivations. Many extensions of CBG(S) are worth looking into in order to move towards real life problems. We intend to look into the following problems:

- Having studied the components in the bargaining agents, we are now in a position to build systems to evolve (as opposed to hard wire) agents for the above bargaining problem, as proposed by Binmore et al [3].
- Each seller (buyer) can negotiate with more than one buyer (seller) at any time, but with limited bandwidth – in this case, it has to decide on when to talk to which buyer (seller).
- Following the above, the seller (buyer) has *n* days to sell/buy *m* pieces of goods – in this case, it has to decide when to stop negotiating with a particular buyer (seller).
- Past experience and market information (such as the average profit and the wealth of each agent) could be included in the model.

Acknowledgements

Tim Gosling is currently sponsored by BT. This work was motivated by, though not supported by this sponsorship. The authors have benefited from discussions with Ken Binmore, Sheri Markose, Abhinay Muthoo and Mathias Kern. This project has also benefited from input by students at University of Essex who have participated in the negotiation game.

Bibliography

1. Axelrod, R., The evolution of strategies in the iterated prisoner's dilemma, in Davis, L. (ed.), Genetic algorithms and simulated annealing, Research notes in AI, Pitman/Morgan Kaufmann, 1987, 32-41
2. Binmore, K., Game theory and the social contract Vol 1: playing fair, MIT Press, 1994
3. Binmore, K., Piccione, M. & Samuelson, L., Evolutionary stability in alternating-offers bargaining games, ELSE working papers 41, ESRC Centre on Economics Learning and Social Evolution, 1996
4. Binmore, K., Game theory and the social contract Vol 2: just playing, MIT Press, 1998
5. Goldberg, D.E., Genetic algorithms in search, optimization, and machine learning, Reading, MA, Addison-Wesley Pub. Co., Inc., 1989
6. Holland, J.H., Adaptation in natural and artificial systems, University of Michigan press, Ann Arbor, MI, 1975
7. Jennings, N.R., On agent-based software engineering, Artificial Intelligence 117 (2), 2000, 277-296
8. Jennings, N.R., Faratin, P., Lomuscio, A.R., Parsons, S., Sierra, C. & Wooldridge, M., Automated negotiation: prospects, methods and challenges, International Journal of Group Decision and Negotiation 10 (2), 2001, 199-215
9. Koza, J.R., Genetic Programming: on the programming of computers by means of natural selection, MIT Press, Cambridge, MA, 1992
10. Muthoo, A., Bargaining theory with applications, Cambridge University Press, 1999
11. Noriega, P. & Sierra, C. (ed.), Agent mediated electronic commerce, Lecture Notes in Artificial Intelligence, Vol.1571, Springer-Verlag, 1999

Table 1: Competition on problems with a wide range of costs and utilities; average profit was 170.8 per game.

Each pair of numbers shows the result of matching one seller to one buyer. For example, when Random Seller played against Random Buyer, the former gained 57 points per game on average (row2, column 2) and the latter gained 114 points per game on average (row 2, column 3). The higher value in each pair is highlighted.

Profits per game	Random Buyer		Simple Buyer		Keen Buyer		Progressive Buyer		Progressive Buyer 2		Easy Buyer		Smart Buyer 1		Avg gain /game
Random Seller	57	114	54	117	39	130	30	133	20	138	55	116	44	122	**49.6**
Keen Seller	99	70	95	75	77	89	38	116	17	120	95	76	51	113	**78.8**
Keen Seller 0	77	93	69	102	49	120	33	122	16	124	87	84	45	120	**62.7**
Keen Seller 2	98	72	92	77	85	82	42	113	18	119	94	77	51	113	**79.9**
Easy Seller	59	111	52	119	41	130	35	136	37	133	60	110	57	114	**56.8**
Smart Seller 1	102	60	78	54	75	56	34	96	39	118	76	64	44	104	**74.6**
Smart Seller 2	101	61	79	54	75	56	40	63	52	80	66	64	53	58	**77.6**
Smart Seller 3	84	83	78	61	73	59	47	78	49	83	77	77	53	62	**76.9**
Smart Seller 4	82	81	100	62	95	60	65	81	32	67	89	81	46	79	**84.8**
Progressive Seller	71	69	69	72	67	70	42	73	28	55	66	75	46	69	**65.0**
Average gain per game		**81.4**		**79.2**		**85.3**		**101.1**		**103.8**		**82.5**		**95.4**	**170.8**

Notes:

1. Each buyer plays against every buyer 10,000 times
2. Both costs and utilities are generated randomly from the range 1 and ***1000***.
3. Both days-to-sell and days-to-buy are randomly generated from the range 3 and 20.
4. In individual games, the cost could be above the utility.
5. All the programs have been implemented in Prolog for its flexibility in symbolic manipulation. Experiments were run on PC under Windows 2000. Random seed 2885 was used in this particular run.

Table 2: Competition on problems with a narrow range of costs and utilities; average profit was 101.0 per game.

Each pair of numbers shows the result of matching one seller to one buyer. For example, when Random Seller played against Random Buyer, the former gained 30 points per game on average (row2, column 2) and the latter gained 53 points per game on average (row 2, column 3). The higher value in each pair is highlighted.

Profits per game	Random Buyer		Simple Buyer		Keen Buyer		Progressive Buyer		Progressive Buyer 2		Easy Buyer		Smart Buyer 1		Avg gain /game
Random Seller	53	48	49	52	30	66	28	63	29	62	49	51	29	60	**44.6**
Keen Seller	70	27	61	34	50	37	23	60	17	54	65	35	23	58	**51.5**
Keen Seller 0	52	48	37	64	30	64	15	67	28	54	56	45	26	64	**40.7**
Keen Seller 2	70	28	60	35	51	36	23	59	17	54	64	36	24	58	**51.4**
Easy Seller	55	46	43	58	33	68	20	81	46	55	53	48	58	42	**51.3**
Smart Seller 1	70	22	49	14	48	14	16	38	31	46	48	20	21	45	**47.3**
Smart Seller 2	70	25	52	14	49	14	14	17	48	10	24	42	13	16	**44.8**
Smart Seller 3	61	35	49	24	48	15	15	23	48	10	23	53	11	17	**42.6**
Smart Seller 4	63	28	66	27	65	18	34	29	15	10	40	61	10	30	**48.7**
Progressive Seller	51	13	43	21	43	12	11	11	12	10	29	35	10	12	**33.2**
Average gain per game		**32.1**		**34.3**		**34.4**		**44.9**		**36.4**		**42.6**		**40.2**	**101.0**

Notes:

1. Each buyer plays against every buyer 10,000 times
2. Costs are generated randomly from the range between 350 and ***500***; utilities are generated randomly from the range between 450 and 500.
3. Both days-to-sell and days-to-buy are randomly generated from the range 3 and 20.
4. In individual games, the cost could be above the utility.
5. All the programs have been implemented in Prolog for its flexibility in symbolic manipulation. Experiments were run on PC under Windows 2000. Random seed 1471 was used in this particular run.

Table 3(a): Analysis of Results in Table 1 – Share of profit by each strategy

Each pair of numbers shows the result of matching one seller to one buyer. For example, on average, when Smart Seller 4 played against Smart Buyer 1, the former gained 27% of the potential profit (46 divided by 170.8 in Table 1) and the latter gained 46% of the potential profit. Together they realised 73% of the potential profit, not 100%, because they did not manage to agree on the price in some games where the profit region is non-empty. The average percentage of profit gained by Keen-Seller-2, against all opponents, say, is 47% (row 5, last column).

% of potential profits gained	Random Buyer		Simple Buyer		Keen Buyer		Progressive Buyer		Progressive Buyer 2		Easy Buyer		Smart Buyer 1		Avg % profit
Random Seller	33%	67%	31%	69%	23%	76%	18%	78%	12%	81%	32%	68%	26%	71%	**29%**
Keen Seller	58%	41%	56%	44%	45%	52%	23%	68%	10%	70%	55%	44%	30%	66%	**46%**
Keen Seller 0	45%	55%	40%	60%	29%	70%	19%	71%	9%	73%	51%	49%	26%	70%	**37%**
Keen Seller 2	57%	42%	54%	45%	50%	48%	25%	66%	10%	70%	55%	45%	30%	66%	**47%**
Easy Seller	35%	65%	31%	69%	24%	76%	20%	80%	22%	78%	35%	65%	33%	67%	**33%**
Smart Seller 1	60%	35%	46%	32%	44%	33%	20%	56%	23%	69%	45%	37%	26%	61%	**44%**
Smart Seller 2	59%	36%	46%	32%	44%	33%	23%	37%	31%	47%	39%	38%	31%	34%	**45%**
Smart Seller 3	49%	49%	46%	35%	43%	34%	28%	46%	28%	49%	45%	45%	31%	36%	**45%**
Smart Seller 4	48%	47%	59%	36%	56%	35%	38%	47%	19%	39%	52%	48%	27%	46%	**50%**
Progressive Seller	42%	40%	40%	42%	39%	41%	25%	43%	16%	32%	39%	44%	27%	40%	**38%**
Average % profit		**48%**		**46%**		**50%**		**59%**		**61%**		**48%**		**56%**	

Table 3(b): Analysis of Results in Table 2 – Share of profit by each strategy

% of potential profits gained	Random Buyer		Simple Buyer		Keen Buyer		Progressive Buyer		Progressive Buyer 2		Easy Buyer		Smart Buyer 1		Avg % profit
Random Seller	53%	47%	48%	52%	30%	65%	28%	63%	29%	61%	49%	51%	29%	60%	**44%**
Keen Seller	69%	27%	61%	34%	49%	37%	22%	59%	17%	53%	64%	35%	23%	57%	**51%**
Keen Seller 0	52%	48%	36%	64%	29%	64%	15%	66%	28%	53%	56%	44%	25%	63%	**40%**
Keen Seller 2	69%	28%	59%	35%	50%	36%	23%	58%	17%	53%	64%	36%	24%	57%	**51%**
Easy Seller	54%	46%	43%	57%	33%	67%	20%	80%	46%	54%	53%	47%	58%	42%	**51%**
Smart Seller 1	69%	22%	48%	14%	48%	14%	16%	38%	31%	45%	47%	20%	21%	45%	**47%**
Smart Seller 2	69%	24%	52%	14%	48%	14%	13%	17%	47%	10%	24%	42%	12%	16%	**44%**
Smart Seller 3	61%	35%	48%	24%	48%	15%	15%	23%	47%	10%	23%	52%	11%	17%	**42%**
Smart Seller 4	62%	28%	65%	27%	64%	18%	33%	29%	15%	9%	39%	61%	10%	30%	**48%**
Progressive Seller	50%	13%	43%	20%	43%	12%	11%	11%	12%	10%	29%	34%	10%	12%	**33%**
Average % profit		**32%**		**34%**		**34%**		**44%**		**36%**		**42%**		**40%**	

Table 4(a): Analysis of Results in Table 1 – Potential Profit Made By Each Pair of Players

Each entry in the table shows the total percentage of potential profits realised by the corresponding players together. For example, Smart-Seller-4 and Smart-Buyer-1 together realised 73% of all the potential profits (row 10, column 8). This is calculated by the total gains by these two players (46 + 79, row 10, columns 14 & 15 of Table 1) divided by 170.8, the average potential profit per game in Table 1. Smart-Seller-4 on average realised 85% of the potential profits with its opponents. On average, the players realised 88% of the potential profits in each game.

	Random Buyer	Simple Buyer	Keen Buyer	Progressive Buyer	Progressive Buyer 2	Easy Buyer	Smart Buyer 1	Average for Sellers
Random Seller	100%	100%	99%	96%	93%	100%	97%	**98%**
Keen Seller	99%	99%	98%	91%	80%	100%	96%	**95%**
Keen Seller 0	100%	100%	99%	91%	82%	100%	97%	**96%**
Keen Seller 2	99%	99%	98%	91%	80%	100%	96%	**95%**
Easy Seller	100%	100%	100%	100%	100%	100%	100%	**100%**
Smart Seller 1	95%	77%	77%	76%	92%	82%	87%	**84%**
Smart Seller 2	95%	78%	77%	60%	77%	76%	65%	**76%**
Smart Seller 3	98%	81%	77%	73%	77%	90%	68%	**81%**
Smart Seller 4	95%	95%	91%	85%	58%	100%	73%	**85%**
Progressive Seller	82%	82%	80%	68%	48%	83%	67%	**73%**
Average for Buyers	**96%**	**91%**	**90%**	**83%**	**79%**	**93%**	**85%**	**88%**

Table 4(b): Analysis of Results in Table 2 – Potential Profit Made By Each Pair of Players

	Random Buyer	Simple Buyer	Keen Buyer	Progressive Buyer	Progressive Buyer 2	Easy Buyer	Smart Buyer 1	Average for Sellers
Random Seller	100%	100%	95%	91%	90%	100%	88%	**95%**
Keen Seller	96%	94%	86%	81%	70%	100%	81%	**87%**
Keen Seller 0	100%	100%	93%	82%	81%	100%	89%	**92%**
Keen Seller 2	96%	94%	86%	81%	70%	100%	81%	**87%**
Easy Seller	100%	100%	100%	100%	100%	100%	100%	**100%**
Smart Seller 1	91%	62%	62%	54%	77%	67%	66%	**68%**
Smart Seller 2	93%	66%	62%	30%	57%	66%	28%	**57%**
Smart Seller 3	95%	72%	63%	38%	57%	75%	28%	**61%**
Smart Seller 4	90%	92%	82%	63%	25%	100%	40%	**70%**
Progressive Seller	63%	63%	54%	22%	22%	63%	22%	**44%**
Average for Buyers	**93%**	**84%**	**78%**	**64%**	**65%**	**87%**	**62%**	**76%**