NORMS FOR ARTIFICIAL AGENT ACTION IN MARKETS

MAGNUS BOMAN, LISA BROUWERS, KARIN HANSSON, CARL-GUSTAF JANSSON, JOHAN KUMMENEJE, HARKO VERHAGEN

Dept of Computer and Systems Sciences,
Stockholm Univ and the Royal Inst of Technology
Electrum 230
SE-164 40 Kista, Sweden
E-mail: mab@dsv.su.se

We summarise our experiences of a number of demonstrators and simulation experiments designed to test the feasibility of using artificial decision making agents in real-time domains, and comment on the significance of our results to autonomous artificial agent action patterns in markets. Some of these experiences are currently being used for the design of an agent-controlled intelligent building. Our main hypothesis is that the use of norms can extend the capability of artificial decision makers beyond what is obtained from implementing individual utility maximizers in keeping with rational choice theory. We also comment on the concept of trust and its relation to norms.

1 Introduction

In many markets, the question whether or not the agents have made optimal deals will remain open after a deal is done, and the agent will not find itself in the same decision situation again. What is important is if the world model of an agent repeatedly facing a certain decision situation could represent this reoccurrence. Unfortunately, agents operating in markets rarely recognise analogous or even identical decision situations. This is problematic, because it is potentially useful to equip the agent with some kind of learning capability so that it may, e.g., recognise a specific bid pattern in a negotiation. Just a slight abstraction on the agent's part will (at least in theory) allow for much better chances of the agent adapting to experiences of its own behaviour and to the feedback it receives. The strong assumptions underlying machine learning mechanisms such as reinforcement learning in agents have limited the applicability of such mechanisms to real-life problems. This has led to other ways of co-ordinating behaviour of market-based agents where the designer is willing to risk having non-deterministic behaviour in the agents. Behaviour not governed directly by the utility functions of the agents is then allowed to emerge. Game theory has been shown to be insufficient for these purposes [1]. Instead, one must turn to general decision theory and in some situations also to more qualitative reasoning methods. We summarise our experiences of a number of demonstrators and simulation experiments designed to test the feasibility of using artificial decision making agents in real-time domains, and comment on the significance of our results to the design of autonomous artificial agent action patterns in markets. Our main hypothesis is that the use of norms can extend the capability of artificial decision makers beyond what is obtained from implementing individual utility maximizers in keeping with rational choice theory.

The next section will present an industrial case used in several of our demonstrators and simulations, viz. agent-assisted choices in the power domain. Section 3 will introduce pronouncers, and section 4 our approach to incorporating norms for multi-agent control. Section 5 comments on real-time issues, while the final section offers conclusions.

2 Control of Multi-Agent Action—The Power Domain

The adoption of utility and probability theory for agent action allows for the monitoring and control of group behaviour, e.g., through incentives, rules, and costs for coalition formation. As an example, assume that a person has access to a small squadron of well equipped agents for finding, negotiating about, and eventually securing electricity contracts. Agent-based decision support is of large interest to consumers as well as producers of electricity: customers want to keep their costs down, while producers have to keep up with new information technology to compete with other utilities. The distributed nature of power distribution itself and the possibility of using the electric grid for agent mobility are two factors which have made agent programming relevant. Moreover, power markets have large price fluctuations due to their dependence on physical factors, such as temperature and weather conditions. These factors make power markets suitable for market-based agents [2].

The person wants to make full use of her squadron, if possible, because her need is too big to be handled by a single contract. She can then let the agents report secured contracts to each other, and let them be aware of the results of each of the other agents in the squadron, in real-time. Specifically, each agent will let its utility function be affected by incoming reports about the results of its fellow agents. In other words, these agents form a coalition, or team, working in real-time. What is individually rational [3] may deviate from what is rational in the coalition, since an agent might, e.g., secure a contract at a good price (i.e., a price having a high utility value for the agent) but without taking into consideration that another agent in the coalition could have got an even better price, e.g., on account of its previous reputation (which could apply to individual agents as well as to coalitions) with the seller. This illustrates several important social aspects, as diverse as the micro/macro link, reputation, trust, the principle of social rationality [4], and real-time group communication protocol design. We view social and rational market-based agents as parts of an artificial social system [5], a multi-agent system (MAS) [6]. The simplest use of agents is when each agent acts as a human assistant, utilised to gain advantages in trading, e.g., because market data, or the market itself, is inaccessible to humans. Already in this scenario, a vast number of research issues must be considered, and one may safely assume that actors on real markets will exploit agent technology to the full, in order to increase efficiency. It is therefore necessary to investigate the possibilities available to them, which include:

- scenario execution
- run-time role switching
- micro-level design of macro behaviour

In scenario execution, the agents are reactive, displaying deterministic group behaviour. An example would be to instruct a set of agents to "start bidding like crazy" for a particular contract until a certain point is reached when all activity stops. This kind of market manipulation is highly sensitive to failure, e.g., a competing agent making a bid over the expectation point early on: there is no graceful degradation to be had if a scenario execution fails.

In most domains, agents benefit from using only some of the alternative behaviours available to them, and each agent can be seen to play a certain role. A simplistic interpretation of the role concept is in co-operating and defecting agents in evolutionary game theory. In real-life domains, however, the set of primitive roles is hard to determine, and it is often equally hard to determine how many agents that should play a certain role. Given that one has succeeded in such a design, one should ask whether it is necessary for an agent to be able to switch role during its operation, or not. For instance, a risk-prone spot market buyer could be forced to become risk-averse if it is running out of capital. A more radical example of a role switch would be between buyer and seller. However, such switches are rare in practice because of the rigidity in the commitment to specific protocols in the agent programming.

Actions are as always governed by constraints. Some of these constraints are operational locally (i.e. pertain to single autonomous agents) while other constraints are global (i.e. pertain to all agents in the MAS, or to a subgroup of agents all belonging to the MAS). An interesting fact is that global constraints can sometimes be replaced entirely by local constraints, making it possible to determine system properties in a truly distributed fashion [7]. In the next section, we will discuss so-called technical norms as constraints.

3 Pronouncers—Decision-Theoretical Studies

The two kinds of decision situations we have studied—choice between markets themselves, and between different offers on a market—can be evaluated as ordinary decisions with risk, given that the agent has some kind of opinion of the uncertainties involved. The agent could build a decision tree in which the choice node in the first kind of choice leads to different markets, and in the second kind of choice to different contract offers. Normative advice could then be given to the agents by a so-called pronouncer, an authoritative entity external to the agents that evaluate decision trees or influence diagrams. The agents call upon the pronouncer when they are faced with a decision situation, after which the pronouncer evaluates the given problem and returns an action to the agent. Normally this would be the action that maximises the expected utility for the agent, but the pronouncer could make use of norms to filter advised actions, in order to account for group utility in a MAS [8], or of constraints representing risk aversion [7]. Each consequence will have a probability and a utility attached. This representation can be interpreted as a formalisation of trust [9], an issue we will return to in the next section. What is initially lacking in the decision trees are the subjective agent assessments regarding probabilities. The utility of each consequence is typically the monetary outcome. Probability assessments are made in view of experience on the agent's part, either through previous decisions made and their outcome, or knowledge about the market acquired as a result of communication with other agents. In the model we have used, any such learning will be restricted to the information relevant for future decisions of the same kind. For example, once an agent has entered a market, communication with other agents on that market will be used to update the probability assessments of the agent regarding markets (not agents or other objects in the domain).

Risk factors affect the price, and yields intricate trade-offs. For instance, if the price in a deal is not fixed, the longer the contract is valid, the more the seller wants, in order to compensate for the risk taken. By contrast, in fixed-price deals, the seller wants the contract to be valid for as long as possible. Trade-offs such as this has led the power companies to insurance companies, and to business strategies in which risks are diminished by insurance and re-insurance. This is especially true for low-probability catastrophic risks [10]. The risk profile will also be used for filtering the normative advice given by the pronouncer in such a way that alternatives deemed too risky will be eliminated from the procedure of rational choice [11]. This will lead, e.g., to risk-averse agents abolishing certain markets.

4 Social Control and Representation of Trust through Technical Norms—Simulation Studies

We have investigated the use of norms as a means to control artificial decision making [8]. Norms are here considered in their technical sense, i.e. as constraints on agent action. For norms to have measurable effect on the operations on a market, the world model of an agent must incorporate data about other agents and their anticipated behaviour on that market. In other words, an open Internet market is unsuitable for implementing norms, and indeed most other forms of social modelling, since each agent repeatedly encounters new actors. By contrast, a market such as Nord Pool (`www.nordpool.no`) can be categorised as semi-open in that the number of participants is in practice limited (by relatively high fees) and fairly static (measured by low entrance and exit rate). This gives an opportunity for normative steering and thus the creation of trust is possible [12]. It is also interesting to note that Nord Pool has been criticised for informal cartel formations and insider transactions [13]. This is a different kind of indication of norms for artificial decision making agents being well worth studying in power markets. The question whether the introduction of agents onto markets similar to Nord Pool would diminish or increase the size of such problems is open [14].

The large number of current studies on trust in electronic commerce has led to a new category of readers of social studies into human trust. However, the entire body of results does not carry over to artificial agents. For instance, Wuthnow identifies three foundations of trust among humans: [15] past performance, resources, and the self of the person being trusted. The last of the three is not applicable to artificial agents. The same is true for the human view, again commented on by Wuthnow, on trust as a measure of optimism. That said, the first two items above naturally apply to artificial agents, whereof the first is obvious, but the second is less studied and more rarely identified. Basically, an agent cannot be trusted if it has insufficient resources for fulfilling its obligations. This correlation between agent co-operation and amount (and quality) of resources available is significant. It suggests that whether or not the resources of the agents in a MAS are made public is important. For instance, a strong incentive for joining a coalition could be that all members openly declare their resources. This can be a first step towards the practical use of the concept of capability [16] in electronic commerce.

Moreover, most research on trust is on trust in institutions, proposing technical solutions such as protocols to assure that agent identity is correctly displayed and interpreted. This is not an issue of trust but an issue of confidence [17]. Trust is instead defined by us (following Gambetta [18], among others) as a form of reliance on other agents, involving beliefs about the likelihood of those agents behaving in a certain way. Note that this definition of trust is closely related to our model of normative decision making[19]: norms form the basis for trust, in the sense that the expectations on the behaviour of other agents is modelled via the norms of the coalition of which they are part. Since membership in a coalition is driven by an expectation of mutual benefit among members, it follows from our definition that a member trusts the members in its coalition more than it trusts other agents. Since the world model of a member typically increases in detail and accuracy with time, agents that remain in a coalition can be said to increasingly trust its members, or else it would have migrated. The agent can also be considered as displaying trust in the coalition itself, i.e. in the norms of the coalition. We have implemented a model for trust in this sense, and tested its efficiency using computer simulations[20]. The simulation model consists of several agents roaming a two-dimensional space. The agents form coalitions having one of the agents act as leader. Every point in the space may contain nothing, one piece of resource 1, one piece of resource 2, or one piece each of resource 1 and resource 2. The agent repeatedly faces the choice to do nothing, to move to another spot, or to take resource 1 or resource 2 (if available). This yields 20 combinations in total.

Each agent has a subjective utility base called a self-model, containing utility values for each of the 20 alternatives. It also has a coalition model containing the utilities the agent presumes the coalition (in the sense of aggregated utility values among members) has for each of the alternatives. The coalition model also expresses the agent's interpretation of the norms that the coalition holds. The degree of autonomy of an agent relative to the coalition determines the extent to which the coalition model is respected when making a decision. An autonomy level of 0.4 expresses that in 40 per cent of the cases, the coalition model is given higher priority than the self-model. Choosing an alternative does not mean a particular event will occur: chance and the skills of the agent also influence the outcome. An agent updates its two models based on the outcome and the result of the feedback that members give in answer to the agent's message reporting on the chosen alternative. The influence of the leader of the coalition also varies. The leadership value codes the weight of the information provided by the coalition leader. A leadership value of 2 expresses that feedback from the leader is twice as important as the feedback of the other members.

The effect of varying leadership value settings and autonomy settings was studied, as was the effect of different types of bootstrapping. The homogeneity of the group of agents was measured in terms of the spread and internalisation of norms. The first factor measures the difference between the norm sets each agent has about the group, while the second measures the difference between the individual norm set and the coalition norm set. A number of hypotheses were tested empirically:

- H1: Variance of behaviour increases with the degree of autonomy.
- H2: Variance of behaviour decreases with the proportional increase of leadership.
- H3: If the self-model is identical to the initial coalition model, agents will adapt their coalition model slower than if they differ (producing higher variance of behaviour).

The norm spreading factor complied with H1 and H2, and with H3, given certain levels of autonomy. The norm internalisation factor did not completely comply with the formulated hypotheses, prompting further inquiry, which we will not report on here (see [19] and [20] for details). Our main point is that the spread and internalisation of norms in coalitions are important measures in empirical studies of artificial social systems.

5 Real-Time Control—The RoboCup Domain

In the power domain, when a buyer agent accesses a possibly restricted fraction of the Internet to buy electricity, it first selects which kind of market to act on. Because of time constraints, it may in reality be unable to access all power markets available to it. The agent may therefore risk going to a market where providers cannot fulfil all the user preferences it embodies, or to a market where the price of the electricity is higher than it would be elsewhere. This means that buying agents will be satisficing bounded rational agents. Some consumers may prefer hydropower to oil power, others may be prepared to pay twice as much for solar power as for fossil fuel generated power. In these cases, the agents will choose which producers to buy electricity from depending on production type. Choices will involve both risk and uncertainty, since each agent will try to fulfil its goal of purchasing electricity given certain constraints (as specified by the user). The time allowed for this purchase varies with the risk-proneness of the user. Some agents may be instructed to secure a contract several days in advance of the start of the delivery, while others might persist in their goal of lowering the price until just an hour remains.

In a successful attempt to prove the feasibility of using pronouncers in real-time domains, we implemented several pronouncers capable of evaluating simple but typical decision trees in less than ten milliseconds [21]. We have used these pronouncers in several applications, including agent-controlled intelligent buildings [22]. Because of a domain requirement of agents being sent on the electric grid, the size of the agent must be kept small, and pronouncers hence became a strong alternative to internal decision modules assisting in agent deliberation. The real-time aspects of our problem have otherwise been studied in the artificial but extremely demanding RoboCup domain [23]. We have had experience of a physical robot team, but our findings related to our simulated team UBU [24] is of more significance to the object of this paper.

Electronic marketplaces share a number of characteristics with RoboCup, e.g., requirements on high dynamicity, concurrency, efficiency, and stability of roles. In RoboCup, there are 22 players and a ball that will typically move with every timestep, and the world changes every 100 milliseconds. We are currently conducting tests of pronouncers in the form of a coach function: an extra player that deliberates under slightly less harsh time constraints than does an ordinary player. A coach can be said to consist of the pronouncer code for decision support, plus code that allows for selections from a library of scenarios. The coach receives the same noisy data from the server about the game at hand as an ordinary player, and it has limited and noisy sensor input, but because it need not move on the field it can concentrate on giving advice to other players. The advice can be proactive, e.g., instructing a player to move its sector of play towards the goal of the opponent in order to allow for more offensive team play. It can also await queries from players that send in decision trees or influence diagrams, and hence act as a calculator. We have recently initiated experiments with augmenting basic calculations of individually rational actions by letting the pronouncer take into account what is socially rational. It is interesting to note that pronouncers used in this way provide an alternative to plan libraries [25], saving agents from dynamic plan revision, which is otherwise a heavy responsibility in real-time systems.

6 Conclusions

The seemingly impossible coupling of decision support and dynamic real-time environments has been realised in a number of demonstrators, as well as for applications involving autonomous agents on semi-open real-time markets. Simulation studies, including the construction and continuing evaluation of a RoboCup team, have provided evidence for norms being an important step when moving from rational choice theory to socially responsible agents. The implementation of pronouncers fast enough for evaluation of choice data in real-time environments has been crucial to our long-term goal of unleashing autonomous artificial decision makers onto real markets.

Acknowledgments

The authors would like to thank Håkan Younes for comments. This research was carried out within the PROMODIS programme (see `www.dsv.su.se/~mab/PROMODIS/Index.html`), sponsored by NUTEK (The Swedish National Board for Industrial and Technical Development).

References

1. Magnus Boman and Love Ekenberg. Decision making agents with relatively unbounded rationality. In Nawarecki, editor, *Proc DIMAS'95*, pages I/28–I/35, 1995.
2. Fredrik Ygge and Hans Akkermans. Power load management as a computational market. In Mario Tokoro, editor, *Proc ICMAS'96*, pages 393–400. AAAI Press, 1996.
3. Colin Camerer. Individual decision making. In J. H. Kagel and A. E. Roth, editors, *Handbook of Experimental Economics*, pages 587–703. Princeton Univ Press, 1995.
4. Nicholas R. Jennings and Jesus R. Campos. Towards a social level characterisation of socially responsible agents. *IEE Proc Software Engineering*, 144(1):11–25, 1997.
5. Rosaria Conte and Cristiano Castelfranchi. *Cognitive and Social Action*. UCL Press, 1995.
6. Magnus Boman and Walter Van de Welde. Multi-agent rationality. In Magnus Boman and Walter Van de Welde, editors, *Multi-Agent Rationality*, number 1237 in LNAI, pages vii–xii. Springer-Verlag, 1997.
7. Love Ekenberg, Mats Danielson, and Magnus Boman. From local assessments to global rationality. *Intelligent Cooperative Information Systems*, 5(2-3):315–331, 1996.
8. Magnus Boman. Norms in artificial decision making. *Artificial Intelligence and Law*, 7:17–35, 1999.
9. Sviatoslav Brainov and Tuomas Sandholm. Contracting with uncertain level of trust. In Cristiano Castelfranchi, editor, *Proc AA WS on Deception, Fraud and Trust in Agent Societies*, pages 29–40, 1999.
10. Love Ekenberg, Magnus Boman, and Joanne Linnerooth-Bayer. General risk constraints. *Journal of Risk Research*, forthcoming.
11. Love Ekenberg, Mats Danielson, and Magnus Boman. Imposing security constraints on agent-based decision support. *Decision Support Systems*, 20(1):3–20, 1997.
12. J. S. Coleman. *Foundations of social theory*. Belknap Press of Harvard Univ Press, 1990.

13. NUTEK. Utvecklingen på elmarknaden: Oktober 1996. Technical report, The Swedish National Board for Industrial and Technical Development, 1997. In Swedish.
14. Tuomas Sandholm and Fredrik Ygge. On the gains and losses of speculation in equilibrium markets. In *Proc IJCAI'97*. AAAI Press, 1997.
15. Robert Wuthnow. The foundations of trust. *Report from the Institute for Philosophy and Public Policy*, 18(3), 1998.
16. Eric Werner. Logical foundations of distributed artificial intelligence. In G.M.P. O'Hare and N.R. Jennings, editors, *Foundations of Distributed Artificial Intelligence*, pages 57–117. John Wiley, 1996.
17. N. Luhmann. *Trust and power*. Wiley, 1979.
18. D. Gambetta, editor. *Trust: Making and breaking cooperative relations*. Basil Blackwell, 1988.
19. Harko Verhagen and Johan Kummeneje. Adjustable autonomy, delegation and distribution of decision making. In *Proc CEEMAS'99*, pages 301–306, 1999.
20. Harko Verhagen. On the learning of norms. In *MAAMAW'99 Poster Proc*. Univ Politecnica de Valencia, 1999. CD-rom.
21. Håkan L. Younes and Magnus Boman. Tools for artificial decision making. In *MAAMAW'99 Poster Proc*. Univ Politecnica de Valencia, 1999. CD-rom.
22. Magnus Boman, Paul Davidsson, and Håkan L. Younes. Artificial decision making under uncertainty in intelligent buildings. In K.B. Laskey and H. Prade, editors, *Proc UAI'99*, pages 65–70, 1999.
23. Hiroaki Kitano, Minoru Asada, Yasuo Kuniyoshi, Itsuki Noda, Eiichi Osawa, and Hitoshi Matsubara. RoboCup: A challenge problem for AI. *AI Magazine*, 18(1):73–85, 1997.
24. Magnus Boman, Jens Andreasen, Mats Danielson, Carl-Gustaf Jansson, Johan Kummeneje, Johan Sikstrom, Harko Verhagen, and Håkan L. Younes. UBU: Pronouncers in RoboCup teams. In H. Kitano, G. Seet, and K. Jagannathan, editors, *Proc RoboCup WS at PRICAI'98*, pages 117–122. RoboCup Federation, 1998.
25. Harko Verhagen and Magnus Boman. Norms can replace plans. In David Kortenkamp, editor, *Proc IJCAI'99 WS on Adjustable Autonomy Systems*, pages 65–67, 1999.