

Revised version appeared in:

Proceedings of the Group Decision and Negotiation Conference

Vol. 2, Recife, 2012, (141-145)

Evaluation of Human-to-Agent Negotiations via an E-negotiation System

Rustam Vahidov & Gregory E. Kersten

InterNeg Research Centre, Concordia University, Montreal, Canada

Abstract

Negotiation is a flexible mechanism for facilitating effective economic exchanges. Electronic negotiations allow participants to negotiate online and use analytical support tools in making their decisions. Software agents offer the possibility of automating negotiation process using these tools. The purpose of this work is to make progress towards outlining design-theoretical principles for agent-enhanced negotiation systems (AENS). This paper describes an electronic marketplace named DIANA (Deal-making system Incorporating Agents in Negotiations and Auctions) that allows involving software agents in negotiations. It also presents the results of experiments in agent-to-human negotiations. Various types of agents have been configured and paired up with human counterparts for negotiating product sale. The paper discusses the results and presents a set of rules for the design of AENS.

© Rustam Vahidov & Gregory E. Kersten

<http://interneg.org/>

Acknowledgments: The publication of the InterNeg Research Papers has been supported by the Natural Sciences and Engineering Research Council, the Social Sciences and Humanities Research Council, and the J. Molson School of Business, Concordia University.

Copyright: The papers' copyright stays with their authors. The paper cannot be reprinted in any form without its authors' explicit consent.

1. Introduction

In a dynamically changing global business environment negotiations represent an important mechanism for facilitating economic transactions. It offers flexibility and active involvement of the participating parties, which other major mechanism categories, such as fixed price, and even auction models, lack. In the course of negotiations parties exchange offers in order to jointly explore the possibilities of finding acceptable solutions. Negotiations involving more than a single (typically price-based) issue allow for more degrees of freedom in search for agreements, which would be beneficial to the negotiators due to the asymmetry of their preference structures. Properly managed and conducted negotiations promise to maximize the mutual benefits of the participants and avoidance of situations characterized as “leaving money on the table”.

Online negotiations supported by electronic negotiation systems (ENSs) allow the parties to exchange offers over the internet (Kersten and Noronha, 1999). In addition to enabling anytime/anywhere mode of interactions, they may also incorporate analytical facilities for supporting negotiators in their preparation and conduct of negotiations in order to achieve the benefits mentioned above. This support can range from such tools as those used for capturing and modeling negotiator’s preferences, to providing active advice and critique, and all the way to complete automation of the negotiation conduct.

Despite the promises and potentials, the existing ENSs have mainly been used in research and educational contexts. Catalogs and auctions have been predominantly employed as means of transacting between businesses and their customers. Could it be that ENSs are not an option for real-world economic parties, and represent, as it were, the “phantom meta-artifacts”? We do not consider this as being the case, and further research is needed into the factors related to ENS adoption and usage (Vahidov, 2012). One possible explanation to the scarcity of real-life ENSs and negotiating websites is that negotiations imply a relatively high cognitive load, especially if multiple issues are involved (e.g. price, warranty, product attributes, shipment, etc.). This load may translate into a prohibitive cost when day-to-day transactions involving people who are not negotiation experts are concerned. Investigating the modes of human-ENS interactions, as well as automation of some structured aspects of negotiations may serve as the key to promoting ENS adoption and usage. Software agents may alleviate the problem of cognitive effort by automating negotiation process while working with customers towards an acceptable deal. Moreover, they can also ensure consistency in reaching negotiation outcomes according to the set policies.

The purpose of this work is to investigate the prospects of negotiations involving humans and software agents in order to make progress towards outlining design-theoretical framework for agent-enabled ENSs. The work represents an early stage towards developing components of design theory for an agent-empowered ENS (AENS). To this end an electronic marketplace system called DIANA (Deal-making system Incorporating Agents in Negotiations and Auctions) has been built. The system was used in experiments with human subjects in order to investigate the effects of various agent strategies on negotiation outcomes, with the purpose of deriving guidelines for AENS design.

2. Related work

Ever since the publication of the seminal paper by Hevner et al. (2004), design-oriented research in IS has attracted a considerable community of followers. Works by numerous various design-oriented researchers have helped to establish the legitimacy of design-type studies and lay the groundwork for theoretical approaches to design science research (e.g. Walls et al., 1992; March and Smith, 1995; Gregor and Jones, 2007). A recent book on the subject explores the similarities between the traditional notion of science and advocates application (with interpretation) of scientific principles to design science research (Vahidov, 2012).

We believe posit that in theoretical approach is required to accumulate and applying design knowledge. In this regard, an ideal formulation of a meta-artifact should be in form of a design theory for that class of artifacts. The concept of a design theory introduced by Walls et al. (1992) includes both type of requirements and type of system solution as key components. The development of these components is guided by the kernel theories. In Gregor & Jones (2007) (2007) formulation one of the key components includes “the principles of form and action”, which could be related to the “type of system solution” mentioned above. These important contributions do not stipulate the exact forms in which different components of a design theory should be formulated structurally and dynamically. Vahidov (2006) has proposed a representational framework for design researcher’s meta-artifacts that includes analytical, synthetic, technological, and implementation layers. Carlsson (2005, p.) has introduced the notion of technological rules to represent design knowledge. The form of such a rule could be expressed as: “If you want to achieve A (outcome) in Situation B (problem) and C (context) then something like action/intervention D can help because E (reason)”. In this work we will present a number of such rules for the design of AENS, which are derived from the results of our experiments. To this end we first need to investigate the related work in electronic negotiations involving software agents.

Research on automated negotiations involving software agents has been extensive (Beam and Segev 1997; Jennings et al. 2001). While thorough coverage of the past work in the area is well beyond the scope of this paper, we will review the representative publications in the context of business exchanges. One could categorize these in accordance with the context of interactions (i.e. C2C, B2B, B2C), and the extent of automation.

One well-known early work in this direction was the construction of the Kasbah electronic marketplace (Chavez et al., 1997; Maes et al., 1999). Targeting primarily the C2C domain the marketplace allowed human users to configure agents, which would then be sent to the marketplace to negotiate with each other. Three types of agents ranging from competitive to the conceding ones were provided. Negotiations included a single issue, i.e. price. In B2B applications software agents have been proposed for automating various aspects of supply chain management. For example, in (Wang et al., 2009) an agent-based architecture has been proposed for dynamic supply chain formation. The agents acting as brokers representing various entities within supply chain negotiated agreements with each other in building up the chain.

There has also been work targeting the B2C transactions. In (Huang et al., 2010) the authors proposed an agent-based architecture for automated negotiations between businesses and consumers. The buyer agents incorporated such components as searcher and negotiator, while seller agents featured negotiator module whose strategy was set by the sales department. It has

been argued by many that complete automation of real-life negotiations, in particular in business contexts does not seem to be a viable solution (e.g. (Lin and Kraus, 2010)). Automation in general is applicable only when tasks concerned are well-structured, which is rarely the case in many business situations. However, since efficient policies can be set for multiple daily interactions with the customers regarding the sales of products and services, it seems that a relatively high level of automation may be feasible.

While the work reviewed above concerns fully automated negotiations, there has been some research into sharing responsibilities between human negotiators and negotiation agents. In (Chen et al., 2005) a system has been proposed where agents actively supported human decision making in the negotiation process. An agent advised the human user on the acceptability of the received offer, helped with the preparation of the counter-offer, and critiqued offers composed by the users when it deemed necessary to intervene. Vahidov (2008) proposes an agent-based architecture was proposed with the purpose of for managing multiple negotiation managements. In this architecture a fleet of agents negotiated deals with customers. These negotiations were monitored by a coordinating agent, which—based on the analysis of situation—instructed the negotiating agents to adjust modify their strategies and adjust reservation levels within the limits of its authority. The overall process was monitored by a human user who could intervene to make changes if necessary.

The current work aims at informing design of electronic agent-populated marketplace and investigation of software agents' performance as compared to human-human dyads while in multi-issue negotiations. Various types of agents following different strategies have been configured for the comparison of their performance.

3. Architecture of DIANA and configurations of agents Subsection

Kersten et al. (2008) proposed a general framework for electronic marketplaces involving humans and software agents called Shaman. The design of DIANA system has been inspired by this framework. Figure 1 shows the simplified architecture of DIANA, focusing on its negotiation support facilities. Negotiation case library stores information on different negotiation cases, including such specifics as the subject of negotiations, issues involved, options for the discrete issues, and other details. The cases can be created by the system administrator. Negotiation engine uses the information from the case library to manage the exchange of offers and counter-offers between the parties. The negotiating parties could be both humans or mixed human-software agent dyads. The analytical toolbox allows modeling of preferences and evaluating received and prepared offers in terms of their overall utilities. These could be used both by humans, as well as agents in the process of exchange. A human user may be the principal of the agent who delegates the task of negotiation to the agent. In this case the principal has to configure the agent to specify its behavior.

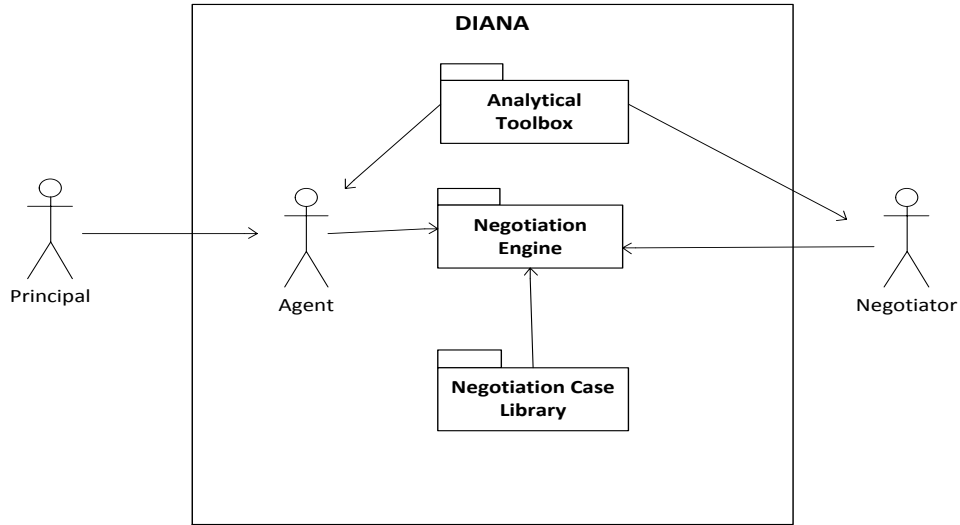


Figure 1. Overall architecture of DIANA

Agents can be configured, in part, by specifying the concession schedule they must follow. In our experiments, we have chosen to use five different concession schedules, three of which were similar to those used in Kasbah experiments (Chavez et al., 1997). These included: competitive, neutral, collaborative, competitive-then-collaborative, and tit-for-tat strategies. The competitive agents (CM) tend to make smaller concessions in terms of utility of generated offers in the beginning of the negotiation period. However, as they approach the end of the period, they would start making larger concessions in search of an agreement (figure 2A).

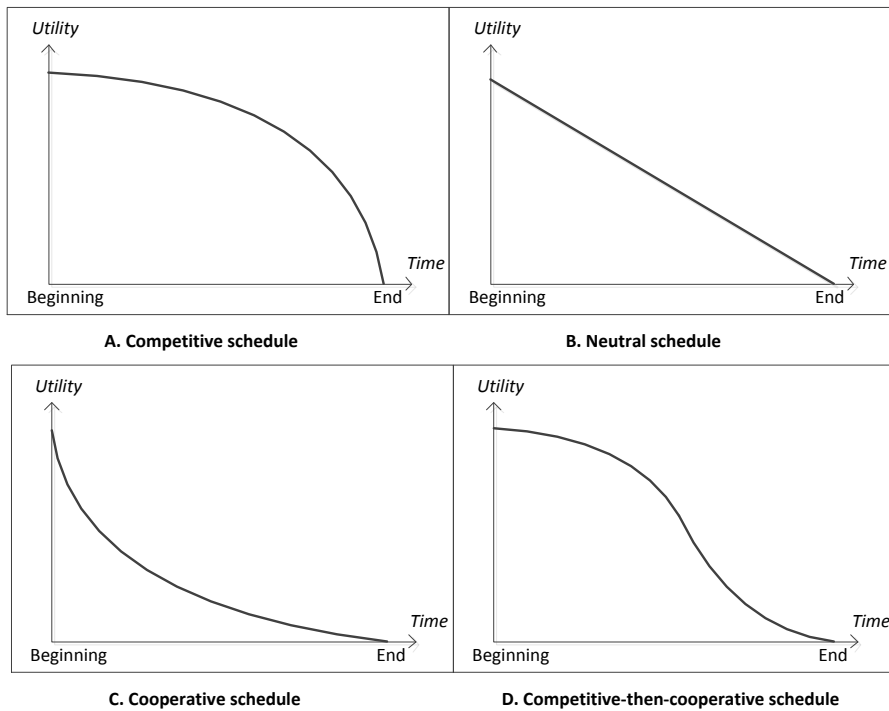


Figure 2. Four schedules

Neutral strategy (NT) dictates that an agent concedes the constant amount of utility regardless of the time period, i.e. the concession schedule is linear (figure 2b). Collaborative schedule (CL) implies making large concessions in the very beginning of the negotiation period in search of a quick agreement. This represents the case where an agent is anxious to sell the product. However, as the agent quickly drops the utility close to the reservation levels, it cannot make large concessions later in the process (figure 2C).

Competitive-then-collaborative schedule (CC) models more complex behavior of the agents. In the beginning of the process an agent behaves competitively, however, in the middle of the negotiation period it changes its profile to a collaborative one. Thus, there is an inflexion point in an agent's schedule (figure 2D).

The reason for introducing this strategy is to imitate the situation when an agent's behavior adjusts due to the overall situation in the market (e.g. the product is not selling well). Moreover, the CC schedule allows introducing less predictable non-obvious behavior, which may be characteristic of human negotiators. (Little circles appearing on the screenshots are used to graphically define the shapes of the curves.)

The final strategy used is tit-for-tat. These agents do not rely on utility calculations. Rather, they watch the opponent moves and simply mirror them in composing counter-offers. In other words, when an opponent makes a new offer an agent determines the difference between this offer and the previous one made by the opponent, and applies the same difference to its own offer. If, say an opponent made a large change to a price, the agent would do the same.

The agent follows the following algorithm. In the beginning of the process it makes an offer that has highest utility to an agent. It then waits for the opponent to respond. If an opponent agrees, the process terminates. If an opponent makes a counter-offer the agent calculates its acceptable utility level according to the concession schedule employed. If the opponent's offer is equal or higher than the acceptable utility, the agent accepts the offer. Otherwise, the agent generates a new offer according to the acceptable utility level. It takes the opponent's offer as a starting point, and employing hill-climbing algorithm changes it to get close to the set utility level. This heuristic method is used instead of analytical one, since most of the issues are not continuous variables. It then sends this offer to the opponent.

4. Experiments

The negotiation case developed for the experimental study concerned the sale of a desktop computer. There were five issues including the price, type of monitor, hard drive, service plan, and software loaded. Each option for each issue had a corresponding level of utility (attractiveness), these levels being different for the buyers vs. sellers. In order to calculate the total utility of the offer the issues were assigned different weights. These were then used in an additive utility function to estimate the level of attractiveness of an offer. Agents used this information in order to decide on the acceptability of the received offers and generate offers.

All agents acted on the seller side, and they were not aware of the buyers' preference structures. The weights were slightly different for sellers than buyers to facilitate tradeoffs, which have been considered one of the key integrative negotiation characteristics (Raiffa et al., 2003). Thus, agents would decide on the utility of the next offer first, according to their concession schedules, and then generate the corresponding offer.

In the current work we were interested in the objective outcomes of agent – human negotiations, as well as subjective variables capturing human perceptions of the process, outcomes and system. The measured variables included the utility of the agreements, and the proportion of agreements achieved. These relate to the economic benefits of agent-human negotiations.

The subjects in the study were university students enrolled in the introductory course on information technology. Thus, the negotiation case was well in line with the learning objectives of the course. The treatments included pairing up the subjects with various types of agents described in an earlier section. We also paired up humans with humans in a control group.

The experiment was conducted via the web, whereby subjects could perform their tasks from any location in an asynchronous mode during a two-day period. The subjects were invited to join the negotiations via email containing the link to the system. Negotiations began by sellers making the first offer. The agent sellers then checked for the status of negotiations at fixed intervals of time (every 3 hours). At those points of time, if they have not received new offers, they would wait until the next period of time elapsed. If an offer was received they would evaluate it and would either accept it, or would make a counter-offer.

Human subjects were free to terminate the negotiation at any time without reaching an agreement with their counter-parts. Upon the completion of the experiment the subjects were asked to answer: “I was negotiating with: 1) a human; 2) a computer; 3) not sure.”

5. Results

For the analysis of the results we have selected only those negotiation instances, which featured at least four offers in total. The rationale for this decision was to include only those cases where the subjects took the task seriously. Thus, we ended up having 436 usable negotiation instances. Of these, 65% ended up in an agreement, while in 35% of cases the agreement was not reached.

Figure 3 shows the results of the question related to whether the participants guessed correctly if they were negotiating with humans or computers. The left side shows the results from human-agent dyads, and the right side shows human-human ones. The leftmost bar in each group indicates the number of responses that read “human”, the middle one relates to “computer” responses, and the last one shows “not sure” responses.

As one can see, the majority of subjects in the agent-human dyads were not sure if they were interacting with the humans or agents (183 responses). This was followed by the group of subjects who had thought they were negotiating with other humans (114). The smallest group consisted of those who guessed correctly that they were interacting with agents (65). It is interesting to note that some subjects in the human-to-human dyads thought they were interacting with a computer (2 out of 30).




Figure 3. “I was negotiating with...” agent - human dyads vs. human - human dyads

The distribution of answers depended on the type of the agent strategy employed. For example, in competitive-then-collaborative category much larger proportion of subjects thought they were negotiating with a human counter-part as compared to those who had an impression they were dealing with a machine (25 vs. 8). This can be explained by the fact that CC concession schedule results in more complex behavior, less obvious behavior that could be more readily ascribed to humans, rather than machines. Similar, though less prominent results were obtained in competitive agent category (33 vs. 15). On the other hand, the collaborative category was the only one where the number of “human” vs. “machine” responses was equal (21 each). Perhaps, the subjects expected their human counterparts to be more competitive, rather than conceding.

Table 1 shows the proportions of agreements for different compositions of dyads. The largest proportion of agreements was reached in the collaborative agent category. This an intuitive result, since collaborative agents make large concessions early in the negotiations process, and thus they have a higher chance of making a deal with the human counterparts. It is interesting to see that human-to-human dyads have a second-lowest record in terms of proportion of agreements made. Thus, the majority of agent-involved dyads have reached more agreements than purely human dyads.

Competitive agents were able to reach an agreement in 53% of cases. Competitive-then-collaborative agents have made agreements in 75% of cases, falling between the CL and CM categories, but higher than neutral category. The lowest number of agreements was achieved in tit-for-tat category. This is the only agent strategy that does not employ utility function, and, thus it does not necessarily drop its utility level to the minimum towards the end of the period. Overall, agent-human pairs achieved agreements in 66% of cases vs. 50% exhibited by HH dyads.

Table 2 compares the utilities of reached agreements for sellers and buyers across different categories. In human-human dyads the sellers achieved much lower utility levels than buyers.

This could be explained by the adopted reference frames. Since both sellers and buyers in this category were undergraduate student subjects, they tended to shift the price levels downwards to what they consider to be acceptable regions. Nonetheless, as it can be seen from the table, the human sellers had reached the lowest levels of utility.

Table 1. Proportions of agreements

Category	Agreements, %
All agent categories	66
Competitive	53
Neutral	70
Collaborative	82
Competitive-collaborative	75
Tit-for-tat	43
Human-human	50

The highest average utility was achieved by tit-for-tat agents (72.4). However, as already mentioned, they performed worst in terms of proportion of agreements reached. In terms of proportion of agreements the competitive agents have performed slightly better than human sellers. However, utility-wise these agents have considerably outperformed their human “colleagues” (63.2 vs. 35.9). Collaborative agents did only slightly better than humans, reaching 36.5 utility. However, they had much higher proportion of agreements. Competitive-then-collaborative agents have reached the average utility level of 40.4, and the neutral ones had a slightly higher value of 43.8. Overall, agents did better than human negotiators (46.8 vs. 35.9).

Table 2. Utilities of agreements

Category	Seller utility	Buyer utility
All agent categories	46.8	65.6
Competitive	63.2	44.9
Neutral	43.8	69.7
Collaborative	36.5	79.0
Competitive-collaborative	40.4	71.9
Tit-for-tat	72.4	36
Human-human	35.9	73.0

In addition to dividing agents into various above configurations we have also had two versions of their algorithms for generating offers. Passive agents generated their offers without taking into account the opponent’s counter-offer, while reactive agents took the opponent’s offer as a starting point and tried to modify it to fit the desired level of utility. As the results indicate, reactive agents were able to achieve 72% agreement rate, while the passive agents only

managed to secure 64% rate. This difference was significant, while there was no significant difference in the utilities of agreements.

6. Design implications

The results show that employing agents in the majority of cases lead to superior results as compared to using human negotiators. The findings also allow us to draw blueprint for the set of design guidelines for the AENS in form of technological rules (Carlsson, 2005). The rules, to remind, have the form “If you want to achieve A (outcome) in Situation B (problem) and C (context) then something like action/intervention D can help because E (reason)”. In our case, the problem on hand is agent – to human negotiation. The reasons for these rules derive from the results of the experiments. Thus, we end up with context/outcome/intervention components. The following text summarizes the empirically supported rules.

1. Overall, taking into account counter-part’s offers while generating counter-offers lead to increased likelihood of agreements. In other words, being adaptive to the opponent pays off.
2. If you want to achieve higher agreement utility, when number of agreements is not of primary concern, use competitive strategies for your agents. For example, in the absence of considerable competition, competitive strategies pay off.
3. If you want to achieve high proportion of agreements, while the utility can be somewhat sacrificed, use collaborative strategies. This could be the case when the competition is high, or the negotiator wants to maintain relationships with the counter-parts, or has excess resources, and under other circumstances when the transaction is very much desirable.
4. If the desired outcome/context may change during the course of negotiations, dynamically adjust the strategies.
5. If it is important that human counter-parts should not guess that they are interacting with a counter-part use complex or dynamic strategies. This rule is supported by the finding that in a relatively complex compete – then collaborate scenario the human counter-parts were least convinced that they are negotiating with a machine.

7. Conclusions

The purpose of this study was to make a progress towards the design an agent-populated marketplace, experimentally investigate the promises of agent-human negotiations in B2C context, and outline rules that could guide the design of AENS. To this end various types of agents were configured to conduct negotiations with human subjects. The question of whether humans were able to tell if they were negotiating with machine has important implications, since if they did they would be, in principle, able to predict the opponents moves. Findings indicate that, in most cases, the subjects were not able to make a correct guess. This is especially true when agents employed a complex concession pattern, i.e. compete-then-collaborate.

In regards with the outcomes the results show that human negotiators performed worst as compared to agents in terms of utility of agreements. They were also second worst in terms of number of agreements. One possibility for future work could be conducting experimental studies where agent and human negotiators could add issues in the course of negotiations.

References

- Beam, C. and A. Segev (1997). Automated Negotiations: A Survey of the State of the Art. *Wirtschaftsinformatik* 39(3): 263-268.
- Carlsson, S. (2005). Developing Information Systems Design Knowledge: A Critical Realist Perspective. *The Electronic Journal of Business Research Methodology* 3(2): 93-102.
- Chavez, A., D. Dreilinger, et al. (1997). A Real-life Experiment in Creating an Agent Marketplace. *Software Agents and Soft Computing*. H. S. Nwana and N. Azarmi, Springer-Verlag: 160-179.
- Chen, E., R. Vahidov, et al. (2005). Agent-supported Negotiations in the E-marketplace. *International Journal of Electronic Business* 3(1): 28-49.
- Gregor, S. and D. Jones (2007). The Anatomy of a Design Theory. *Journal of the Association for Information Systems* 8(5): 312-335.
- Hevner, A., S. T. March, et al. (2004). Design Science in Information Systems Research. *MIS Quarterly* 28(1): 75-105.
- Huang, C.-C., W.-Y. Liang, et al. (2010). The Agent-based Negotiation Process for B2C E-commerce. *Expert Systems with Applications* 37(1): 348-359.
- Jennings, N. R., P. Faratin, et al. (2001). Automated Negotiation: Prospects, Methods and Challenges. *Group Decision and Negotiation* 10(2): 199-215.
- Kersten, G. and S. J. Noronha (1999). WWW-based Negotiation Support: Design, Implementation, and Use. *Decision Support Systems* 25: 135-154.
- Kersten, G. E., R. Kowalczyk, et al. (2008). Shaman: Software and Human Agents in Multi attribute Auctions and Negotiations. *Negotiation, Auctions, and Market Engineering*. H. Gimpel, N. R. Jennings, A. Ockenfels and C. Weinhardt, Springer Berlin Heidelberg. 2: 116-149.
- Lin, R. and S. Kraus (2010). Can Automated Agents Proficiently Negotiate with Humans? *Communications of the ACM* 53(1): 78-88.
- Maes, P., R. H. Guttman, et al. (1999). Agents that Buy and Sell. *Communications of the ACM* 42(3): 81-87.
- March, S. T. and G. F. Smith (1995). Design and Natural Science Research on Information Technology. *Decision Support Systems* 15: 251-266.
- Raiffa, H., J. Richardson, et al. (2003). *Negotiation Analysis. The Science and Art of Collaborative Decision Making*, Cambridge, Harvard University Press.
- Vahidov, R. (2006). Design Researcher's IS Artifact: a Representational Framework. 1st International Conference on Design Science Research in Information Systems and Technology, DESRIST'2006 Claremont, CA.
- Vahidov, R. (2008). Situated Decision Support Approach for Managing Multiple Negotiations. *Negotiation, Auctions, and Market Engineering*. H. Gimpel, N. R. Jennings, G. E. Kersten, A. Ockenfels and C. Weinhardt, Springer Berlin Heidelberg. 2: 179-189.
- Vahidov, R. (2012). *Design-type Research in Information Systems: Findings and Practices* Hershey, PA IGI Global.
- Walls, J. G., G. R. Widmeyer, et al. (1992). Building an Information System Design Theory for Vigilant EIS. *Information Systems Research* 3(1): 36-59.
- Wang, M., H. Wang, et al. (2009). Agent-based Negotiation and Decision Making for Dynamic Supply Chain Formation. *Engineering Applications of Artificial Intelligence* 22(7): 1046-1055.