

# A classification scheme for negotiation in electronic commerce\*

Alessio R. Lomuscio<sup>†</sup>, Michael Wooldridge\* and Nicholas R. Jennings<sup>‡</sup>

<sup>†</sup> Department of Computing  
Imperial College of Science, Technology and Medicine  
London SW7 2BZ, United Kingdom  
A.Lomuscio@doc.ic.ac.uk

\* Department of Computer Science  
University of Liverpool  
Liverpool L69 7ZF, United Kingdom  
M.J.Wooldridge@csc.liv.ac.uk

<sup>‡</sup> Department of Electronics and Computer Science  
University of Southampton  
Southampton SO17 1BJ, United Kingdom  
nrj@ecs.soton.ac.uk

March 31, 2000

## Abstract

In the last few years we have witnessed a surge of business-to-consumer and business-to-business commerce operated on the Internet. However many of these systems are often nothing more than electronic catalogues on which the user can choose a product which is made available for a fixed price. This modus operandi is clearly failing to exploit the full potential of electronic commerce. Against this background, we argue here that in the next few years we will see a new generation of systems emerge, based on *automatic negotiation*. In this paper we identify the main parameters on which any automatic negotiation depends. This classification schema is then used to categorise the subsequent papers in this book that focus on automatic negotiation.

---

\*This research was supported by the EPSRC under grant GR/M07076.

# 1 Introduction

The term “electronic commerce” (e-commerce) generally denotes an advanced step of modern commerce in which the figures of buyer and seller are replaced by electronic entities [BB97]. It is widely believed that e-commerce will reduce costs for enterprises and provide customers with better bargaining tools. For example, automation can exploit short-term contracts in combinatorially complex settings [San99] and allow extensive comparisons to be performed between a wide variety of goods. Recent estimates put revenues of e-commerce at US\$95 billion for the year 1999 with a forecast of US\$1.3 trillion by 2003 [Res99].

Given the nature of the tasks involved in e-commerce, it has been suggested that agent technology [WJ95] will play an important role in its development (see [GMM98] for a survey). In such settings, agents will play the roles of buyer, seller, mediator [RNSP97], facilitator, and information provider. In so doing, agents will automate part or all the business tasks involved in e-commerce.

Most current e-commerce applications (servers and clients) may be classified as *first-generation* systems. That is, servers are connected to the Internet and they let the user browse through catalogues containing well-defined commodities (e.g., flights, books, compact discs, computer components, etc.) and make purchases typically by means of a credit card transaction. In this case the system’s client can be any basic Internet browser and therefore the users’ actions, being reduced to select/accept choices, are somewhat limited. However several more advanced forms of first generation clients, such as Jango [JAN] and Bargainfinder [BF], have also been introduced. The main feature of these clients, frequently called *shopping assistants*, is the ability to perform *merchant brokering* on behalf of the user. Merchant brokering is the stage in the consumer buying behaviour where the buyer has chosen the product and is looking for the best deal to purchase it<sup>1</sup>.

Other shopping assistants (e.g., PersonaLogic [Per], Firefly [FF]) offer their services earlier in the transaction by helping the buyer in the *product brokering* stage. This is the stage in the consumer buying behaviour where a decision is made about what product the buyer needs. Yet more advanced shopping agents, notably Tete-a-Tete [TAT], attempt to follow the buyer in more than one phase of the transaction by performing both product broker-

---

<sup>1</sup>The six main stages of consumer buying behaviour can be broadly thought of as: need identification, product brokering, merchant brokering, negotiation, payment and delivery, and service and evaluation. See [Nic66, HS69, EB82, Bet79] for more details on consumer buying behaviour.

ing, merchant brokering, and even some primitive forms of negotiation.

A key problem with all of these first-generation systems, however, is that they are too focused on one aspect of the transaction: price. Although price is clearly important, it is very often the case that neither customers nor enterprises have price as their only concern. Enterprises usually aim to acquire or maintain regular customers, whilst these customers are, in turn, often interested in non-price aspects of the purchase (such as warranty, delivery time, and past-record of the company). To cope with this need, more advanced forms of shopping assistant are required. Such software will have to deal with (interact over) the complete range of issues that are relevant to their user.

When faced with the need to reach agreement on a variety of issues, humans make use of *negotiation*. Similarly, we believe that automated negotiation will become the dominant mode of operation for shopping assistant agents. This *automation* of negotiation can significantly reduce negotiation time (making large volumes of transactions possible in small amounts of time) and can also remove some of the reticence of humans to engage in negotiation (e.g., because of embarrassment, personality, etc.). For these reasons, the formalisation of negotiation has received a great deal of attention from the multi-agent systems community throughout the past two decades [DS83, KL89, KM90, RZ94]. Through such endeavours, it is becoming possible to develop practical automated negotiators. This will, in turn, facilitate the development of fully-fledged shopping assistants that are able to perform negotiation on behalf of users.

Negotiation strategies and their corresponding properties depend heavily on the specific characteristics of the scenarios under consideration. For example, economists and game theorists distinguish between scenarios in which the value of a good is common to all the agents (for example a typical shrink-wrapped good such as a CD), and those in which it is a private characteristic, which differs from agent to agent (such as a painting, which different agents may value differently) [OR94]. Whether a good has a common value, or whether its value differs from agent to agent, the valuations themselves may be common knowledge or private. These are but two illustrations of parameters upon which any negotiation mechanism crucially depends. However, as we shall see, many other parameters also play a key role in shaping the negotiation. Against this background, the principle aim of this paper is to define (for the first time) the *negotiation space* for electronic commerce. That is, we aim to identify the possible parameters that can be used to classify any negotiation mechanism for electronic commerce. Such a classification is an important step for the development of more so-

phisticated shopping assistants because it *defines* and *delimits* the design space for agent interactions. Thus, the classification in this paper should be seen from an analytical point of view. It does not offer any immediate practical application, but, instead, it aims at providing a conceptual framework on which protocols and strategies for negotiation can be classified and reasoned about. We argue that the components of the negotiation space that we identify constitute a complex space of possible games; indeed, we have so far been unable to find two distinct negotiation scenarios that would be described in the same manner in the classification. We illustrate our taxonomy by classifying the other papers in this volume that deal with negotiation.

The remainder of this paper is structured as follows. In section 2 we discuss and classify the main variables upon which a negotiation for electronic commerce depends. In section 3 we classify the papers in this volume according to our taxonomy. We conclude in section 4 with some discussion on the future direction of automated negotiation in e-commerce.

## 2 Negotiation space

In its broadest sense, automated negotiation involves the design of the high-level protocols for agent interaction. The theme of negotiation is present in many different fields and, as a result, several definitions have been proposed in the literature [RZ94, Lee96]. It is not within the scope of this paper to discuss the merits of these definitions, and so we simply adopt the following definition:

Negotiation is the process by which group of agents communicate with one another to try and come to a mutually acceptable agreement on some matter.

Two basic components are important when designing an automatic negotiation system: the *negotiation protocol* and the *negotiation strategies*. The former specifies the “rules of encounter” between the negotiation participants. That is, the protocol defines the circumstances under which the interaction between the agents takes place: what deals can be made and what sequences of offers are allowed. In general, agents must reach agreement on the negotiation protocol to use before negotiation proper begins. Reaching agreement on a negotiation strategy may itself be done by *meta-level* negotiation, but we do not consider this process here.

An agent’s negotiation *strategy* is the specification of the sequence of actions (usually offers or responses) the agent plans to make during the ne-

gotiation. There will usually be very many strategies that are compatible with a particular protocol, each of which may produce a very different outcome. For example, an agent could concede at the first round or bargain very hard throughout the negotiation until its private timeout is reached. It follows that the negotiation strategy that an agent employs is crucial with respect to the outcome of the negotiation. It should also be clear that strategies which perform well with certain protocols will not necessarily do so with others: the choice of strategy to use is thus a function not just of the specifics of the negotiation scenario, but also of the protocol in use.

A *negotiation mechanism* consists of a negotiation protocol together with the negotiation strategies for the agents involved. There are some properties that are generally considered desirable for a negotiation mechanism (cf. [San99, pp12–14] and [RZ94, pp20–22]).

**Computational efficiency:** Ideally we seek a negotiation mechanism that is computationally efficient. Although users constantly have at their disposal faster machines, there will be little scope for negotiation algorithms that are, say, EXPTIME-complete<sup>2</sup>.

**Communication efficiency:** All things being equal, we would rather have a mechanism that handles communication among the agents in an efficient way. Broadcasting to all the agents in the system, for example, may not be ideal with this respect.

**Individual rationality:** A mechanism should be individually rational for all the agents involved. In other words, it should be in an agent's independent interest to participate in negotiation — otherwise, a rational agent will not. If considerations of group utility need to be taken into account, they can be made a component of each agent's private utility.

**Distribution of computation:** Mechanisms that distribute the computation over the agents involved are preferable to ones in which one server is performing all the computation for the whole system. This is preferred for many reasons, including the desire to avoid the disruptive effects of a single point of failure, and performance bottlenecks.

---

<sup>2</sup>One should be careful not to confuse the computational complexity associated with *designing* a negotiation strategy compared to the complexity of *executing* a strategy. The former can usefully be carried out offline beforehand, thereby reducing the computational costs of the agents. The latter obviously needs to be carried out at run-time and must therefore be manageable.

**Pareto efficiency:** An outcome is Pareto efficient if there is no other outcome that improves the lot of one agent without making another agent worse off. All other things being equal, Pareto efficient solutions are preferred over those that are not.

**Symmetry:** Generally in e-commerce settings, no one agent has complete control of the game. Therefore we prefer mechanisms that guarantee symmetry in terms of power of the agents.

Even with these broad guidelines in mind, many different mechanisms can be designed. These mechanisms very much depend on the specific characteristics of the negotiation space, i.e., the set of all possible negotiation games. In the next section we investigate this space in more detail in order to try and separate out the key orthogonal components.

## 2.1 Parameters of the negotiation space

As already noted, the design of an appropriate negotiation scenario crucially depends on a number of parameters, which can be seen to generate a space of possible negotiation scenarios. A preliminary attempt to classify this space has already been proposed by Wurman *et al.* [WWW98]. In particular, they identify variables upon which auction servers can be designed. Some of the parameters introduced there are also present in our classification; but since we are not only concerned with auctions, our classification is more general.

We divide the parameters on which the negotiation can take place into the following broad characteristics:

- cardinality of the negotiation;
- agent characteristics;
- environment and goods characteristics;
- event parameters;
- information parameters; and
- allocation parameters.

Under this characterisation, an agent's negotiation strategy is simply one of its characteristics. We structure the characteristics of the protocol by analysing in more detail some key parameters such as events, information parameters, and allocation. The protocol as a whole results from the definition of all these.

### 2.1.1 Cardinality of the negotiation

We can distinguish between the cardinalities of the negotiation domain itself, and of the interactions that take place, as follows:

- negotiation domain: single-issue or multiple-issue; and
- interactions: one-to-one, many-to-one, many-to-many.

The domain of negotiation can be thought of as simply the set of tuples over which the agents negotiate. These tuples represent the issues on which an agreement is to be made. Elements of these tuples correspond to issues such as price, quality, warranties, delivery time, and so on (see for example [FSJ98]). In the case of single-issue negotiation (for example where the only issue is price) the tuples are singletons. In the case of multiple-issue negotiation, the different issues might be related by some publicly agreed utility function. This makes the process of bid formulation easier because each agent is then able to compute the utility function of each agent involved in the negotiation [VJ00].

Interactions between agents can be classified in terms of the number of agents participating in the negotiation. *One-to-one* negotiation (where one agent is negotiating with exactly one other agent) is important for both theoretical and practical reasons. Theoretically, it is important because of the technical difficulties that this apparently simple setting provides; for example most games that result from one-to-one negotiation can be proven to have multiple equilibria [OR94], and a naive application of game-theoretical tools is therefore not possible. Practically, it is important because of the emerging role of one-to-one relations associated with business-to-business e-commerce scenarios.

*Many-to-one* negotiation (where many agents negotiate with just one agent) is the standard setting of auctions, which have been popular on the Internet for some time now [QXL, EBA]. In this setting one agent plays the role of the seller, while the many play the role of the buyers.

Finally, *many-to-many* negotiation (where many agents negotiate with many other agents) constitutes the most complex scenario. The continuous double-auction is the most complex of these scenarios that is actually in use [FR93].

Note that whether a given negotiation is one-to-one, one-to-many, or many-to-many is not determined simply by the number of agents. For example, it is conceivable to have a large set of agents, but arrange them such that any negotiation takes place on a one-to-one basis.

### 2.1.2 Agents characteristics

Agents are the computational entities that participate in the negotiation process. Each agent is assumed to be capable of rating its preferences, so that it can evaluate and choose between different deals. We can further characterise agents depending on their:

- role;
- rationality;
- knowledge;
- commitment;
- social behaviour; and
- bidding strategy.

Agents can interact in the negotiation as buyers, sellers, or both (as in a continuous double-action). This is their type: the role they play in the negotiation. For most negotiation domains, buyers and sellers are obviously the leading figures, but in other scenarios, (auctions for example), intermediaries can have an important role.

Rationality can be *perfect* or *bounded* [OR94]. The assumption of perfect rationality generally amounts to the agents being able to perform arbitrarily large computations in constant time. In all practical scenarios agents do not have the computational power (or the ability) to perform such calculations, and they are forced to bid or withdraw on the basis of finite computations. Thus, negotiation models that assume perfect rationality (e.g., many of those coming from game theory) have to use approximations in practice, whereas models that explicitly assume bounded rationality are more realistic in this sense.

Agents have knowledge about the goods they bid for, and possibly some knowledge about how other agents value the same and other goods. Depending on how such knowledge is distributed, agents may choose different bidding strategies. Whether or not agents hold private information is crucially important to design the agent's bidding strategy. For example, internal deadlines, and the valuation of the opponent's utility functions can be important parameters of an agent strategy.

Various levels of commitment can be present. For example, after having made an offer, agents might be obliged to stop bidding for similar goods



until an acceptance or counter offer is received. Alternatively, agents can have the mechanism and process for reneging upon contracts built into the original negotiation (e.g., [San99, FSJ98]).

Agents can act as individually self-interested entities, as altruistic units of a society, or they can strike a balance somewhere in between. Furthermore, whatever their attitude, they can play as distinct entities or effectively pursue team-formation in order to get better deals when convenient [VJ00]. Another possibility is for agents to form coalitions — such coalitions can be a powerful tool for obtaining better deals in some domains.

An agent's bidding strategy is ultimately the component that decides about placing or accepting offers, making counter offers or withdrawing from negotiation. Although conceptually independent from the other parameters, it is reasonable to imagine the bidding strategy to be somehow related to the commitment, the knowledge, the rationality, and the social behaviour of an agent [FSJ98]. A standard assumption is that agents are individually rational, in that their bidding strategy would be against accepting deals that result in making them worse off.

### 2.1.3 Environments and goods characteristics

The negotiation environment can either be static or dynamic. A static environment is one whose variables (e.g., prices of important commodities) are constant over time; a dynamic environment is one in which these change over time. Clearly, there are various degrees of how static or dynamic environments can be, and this is only a simplification. Even so, the dynamicity of the environment might affect the design of the utility function of the agents in a subtle way. The utility functions of the agents reflect their preferences; so, while in a static environment one can imagine an agent that does not learn during the process and maintains a fixed utility function, this behaviour would be less likely to produce a positive payoff in a very dynamic environment.

The characteristics of the goods also crucially define the negotiation protocol:

- private/public value of the goods; and
- nature of the goods.

Agents can value the good differently depending on whether it is intended for private use (e.g., a cake) or whether its value depends on how the other agents value it (e.g., bonds). Frequently, both private and public valuations

play a part. For example, when buying a car one has to consider both one's own preferences and how the car will preserve its value over time in case one should be interested in selling it.

The object of the negotiation can either be a discrete or continuous set (of goods). For example, a negotiation over a deadline (delivery dates, completion of contracts, etc.) is a negotiation over an uncountable set because time is intrinsically continuous. Usually, a simplification is made with respect to the granularity of the domain, and countable or even finite sets are used. For example, with a deadline one could use, days, hours, or minutes, depending on how fine the model needs to be. This simplification considerably reduces the size of the negotiation space. The simplest case arises when the set is a singleton.

#### 2.1.4 Events parameters

The negotiation protocol is mainly influenced by the ways in which the offers and other events that take place during the negotiation are regulated. Indeed, what follows forms an important part of the specification of the protocol of the negotiation. We can distinguish between:

- bid validity;
- bid visibility;
- clearing schedule and timeouts; and
- quotes schedule.

The first item specifies an important part of the protocol: the criteria for validity of the bids. To be valid, bids often have to be offered at an appropriate time and must satisfy some constraints on their value. For example, in an English auction bids can be made when the auctioneer is calling for bids and must be progressively higher in value. Similarly, procedures for placing bids might be present during a negotiation.

The visibility of the bids is only relevant in the case of many-to-one or many-to-many negotiations. At the extremes, bids can be private messages passed between buyer and seller, or broadcast to all agents. Alternatively, we can have configurations in which only subsets of agents see some selected messages (especially useful if coalition formation is permissible).

A “clear” is the event producing a (temporary) allocation between buyer and seller. Clears can be scheduled at random times or following some other events (such as the first offer matching the buyer's request). For example,

during the bidding phase of an English auction each round terminates with a temporary allocation of the good being auctioned to the prospective buyer that meets the auctioneer’s call.

Timeouts determine the closing of the negotiation, therefore they transform clears into “final clears”, i.e., a final agreement between buyer and seller about the transaction. Clears and timeouts also depend on the allocation parameters (see below).

If third-party quotes are generated during the negotiation process, their number and frequency are also part of the events and they need to be regulated. Intuitively, an excessive number of request for quotes can significantly slow down the negotiation mechanism.

### 2.1.5 Information parameters

Both before and during negotiation, information other than bids may pass between the negotiation participants. These messages are either information that can help buyers and sellers reach agreements, or information that can help limit the noise produced by agents trying to buy and sell goods. Such messages can be beneficial in order to save computational time of the agents.

Among the many possible useful messages, we can distinguish between:

- price quotes; and
- transaction history.

Quotes generated by prospective buyers requesting an indicative price from a seller before starting a negotiation can be useful to all parties, as they can reduce negotiation time.

The history of similar transactions can also be requested, or unilaterally provided, by a seller agent in order to give credibility to the information it is offering to potential buyer agents. It can be guaranteed by trusted third-parties. Together with third-party quotes, transaction histories can form the basis for argumentation-based negotiation. Such protocols [SJNP98, KSE98] aim to bridge the gap between how negotiation is performed in human and artificial societies. In the former, expert human negotiators often focus on the reasons why an offer is not acceptable and try and persuade their counterpart of the characteristics that an agreement will have to include. In the latter, artificial negotiators traditionally are only able to propose offers to the counterpart without being explicitly able to motivate an agreement. This is often seen as a severe limitation and can limit the flexibility of the negotiation.

### 2.1.6 Allocation parameters

The allocation parameters only apply in many-to-one and many-to-many scenarios. They govern the winner of an auction when more than one agent has shown an interest in the good.

Allocation is studied in auction theory [Vic61]. The  $M$ th and  $(M + 1)$ -th price allocation policies cover most scenarios, where  $M$  is the number of received bids. Note that if the negotiation is multi-issue, in order for this mechanism to work there has to be a commonly agreed function that “weights” the different issues producing a value of utility. Without this, it would not be possible to order the offers as required by the allocation policy.

## 3 Papers in this volume

This section briefly reviews the papers contained in this volume in terms of the conceptual schema we have developed.

### Paurobally and Cunningham [PC00]

Paurobally and Cunningham [PC00] present a formal model of e-commerce systems in which a phase of automatic negotiation is present. The model uses the machinery of dynamic logic [Gol92] to represent the states and processes of a negotiation system.

Dynamic logic has the advantage of providing a precise and unambiguous language for representing and reasoning about actions and states. For this reason, the framework presented in [PC00] is very general and can be applied to a variety of different settings. The contribution of the paper, therefore, does not lie in a new negotiation strategy for a certain protocol, but, rather, it provides a language for performing inconsistency checks in the specification and in investigating issues such as concurrency.

Nevertheless, the framework does commit itself to a particular subset of all the negotiation games: it deals with one-to-one negotiation and it focuses on a single issue of negotiation.

### Vetter and Pitsch [VP00]

Vetter and Pitsch [VP00] present CASBA [CAS], a project devoted to designing and implementing a flexible and highly automated electronic marketplace. Technically the CASBA system consists of a server and three applets, all written in JAVA. The client-server communication is implemented in CORBA and XML.

Vetter and Pitsch recognise that current e-commerce systems do not address all the stages of consumer buying behaviour (see above), and their proposal should be seen as a preliminary attempt to bridge this gap. As previously mentioned, one of the stages of consumer buying behaviour that is not usually covered is that of negotiation.

- The negotiation protocol described in [VP00] is one of alternating offers. CASBA allows full parallel many-to-many negotiations over multiple attributes.
- Agents' roles are defined in advance (buyers or sellers) and the agents are assumed to be able to evaluate an offer by computing its value in terms of a given (private) utility function. Agents have private information in the sense that their utility function is not necessarily known to other negotiation participants. Agents' commitment is total; the protocol guarantees that a deal that has been agreed upon by two parties can no longer be negotiated upon. A feature of the CASBA system is that the agents' bidding strategy can be programmed via an expert-system shell, or chosen in advance from a set of pre-programmed strategies.
- The environment is static; indeed, the utility functions are not assumed to change during the negotiation process. This reflects the fact that the value of the attributes of the goods being traded is constant. The value of the goods, which are multi-attribute, is private.
- No timeouts or requests for information are present and bids are private.

### **Oliveira and Rocha [OR00]**

Oliveira and Rocha [OR00] present an e-commerce architecture in which special attention is reserved for the negotiation phase. The architecture focuses on a multi-agent system in which a designed agent, called *market agent*, plays the role of the market coordinator. The other entities are either buyers or sellers.

What is particularly interesting in this work is that the decision making apparatus of the agents is built on a reinforcement learning algorithm. This approach builds upon previous work by the authors [OFSG99, OFJ99] in which Q-Learning algorithms [SB98] are developed and tested in market

scenarios. The main contribution of [OR00] is to lay the foundation for a Q-Learning mechanism suitable for multi-attribute negotiations in e-commerce settings.

The protocol of negotiation can succinctly be described as turn-based alternating offer, in which the market agent initiates the negotiation by announcing the need for a product together with a deadline. In response to this announcement, agents can submit bids to the market agent. Bids are computed by the agent's reinforcement learning mechanism, and are private. Upon receipt of the bids, the market mechanism computes the most valuable and sends feedback to the agents on how to improve their bids. This process continues until the deadline is reached or a bid has been judged as acceptable.

In terms of our schema:

- The negotiation domain is multiple, as any offer consists of a variety of issues. Agents interact on a many-to-one basis. One agent is offering a good, through the market agent, and the others bid for it.
- The role of the agents can be of three different types: buyer, seller, or market agent. Agents have bounded rationality and they have private information (each agent's utility function is private). The agents are self-interested and bid according to a Q-Learning algorithm described in the paper. The commitment of the agents to follow the agreement is ensured by the protocol.
- The environment is constant, i.e., the value of the goods does not depend on external time dependent factors.
- Goods have private value. They can be discrete or continuous in nature. In general, they have more than a single attribute.
- The validity of the bids produced by the agents is guaranteed by the market agent. When bids are put forward, they are visible only to the market agent. Timeouts for completion of a negotiation are public and announced by the market agent at the beginning of the negotiation. Clears are produced and sent to the appropriate agent.

## 4 Conclusions

The last few years have witnessed an unprecedented expansion of business carried out online. While these systems perform automatic transactions,

thereby reducing costs, it is often the case that they amount to little more than electronic catalogues on which credit card payments can be arranged online. However a new, more flexible, class of e-commerce systems are starting to emerge. In such systems, automated software agents participate in trading activities on behalf of their user. A key aspect of these systems is the richness and flexibility of the inter-agent interactions. Here we focused specifically on automated negotiation since we consider this to be the key type of interaction in such systems. Automated negotiation, like its human analogue, is a very broad and encompassing discipline. For this reason, it is important to understand the dimensions and range of options that are available. As a preliminary step in this direction, we have developed a classification scheme that is specifically targeted at automated negotiation for e-commerce settings. This classification scheme was illustrated on the papers contained in this volume and the results are summarised in Table 1.

Having specified the negotiation design space, the next step is to identify the tools and techniques that are appropriate for specific regions of this space. Thus, we believe that no one technique or method of approach will come to dominate in this field [JPSF00]. Rather, there will be a range of alternatives that may be selected according to domain circumstances. The act of mapping the negotiation space in this manner will produce a design repository for interactions in agent-mediated electronic commerce. This repository will then constitute a valuable resource that will enable design expertise and know-how to be shared between developers.

<i>Work</i>	<i>Cardinality</i>	<i>Agent Characteristics</i>	<i>Goods</i>	<i>Events parameters</i>	<i>Info parameters</i>
[PC00]	Negotiation issues: 1 Interactions: one-to-one	Agent roles: Buyer/Seller Bidding strategy: N/A Rationality: N/A Commitment: N/A Behaviour: N/A	Value: N/A Nature: N/A	Bids validity: N/A Bidding: N/A Clearing: N/A Quotes-sched: N/A	Price-quotes: N/A trans hist: N/A
[VP00]	Negotiation issues: Multi-issue Interactions: many-to-many	Agent roles: Buyer/Seller Bidding strategy: via an expert system Rationality: Bounded Commitment: total Behaviour: self-interested	Value: Private Nature: discrete	Bids validity: N/A Bidding: private Clearing: N/A Quotes-sched: N/A	Price-quotes: no trans hist: no
[OR00]	Negotiation issues: Multi-issue Interactions: many-to-one	Agent roles: Buyer/Seller/Market-Agent Bidding strategy: Q-Learning Rationality: Bounded Commitment: total Behaviour: self-interested	Value: Private Nature: discrete	Bids validity: N/A Bidding: private Clearing: Yes Quotes-sched: N/A	Price-quotes: no trans hist: no

Table 1: A table summarising the works analysed in section 3



## References

- [BB97] J. P. Bailey and Y. Bakos. An exploratory study of the emerging role of electronic intermediaries. *International Journal of Electronic Commerce*, 1(3):7–20, 1997.
- [Bet79] J. Bettman. *An information processing theory to consumer choice*. Addison-Wesley, 1979.
- [BF] <http://bf.cstar.ac.com/>.
- [CAS] <http://www.casba-market.org/>.
- [DS83] R. Davies and R. G. Smith. Negotiation as a metaphor for distributed problem solving. *Artificial Intelligence*, 20(1):63–109, 1983.
- [EB82] J. Engel and R. Blackwell. *Consumer behavior*. PCBS college publishing, 4th edition, 1982.
- [EBA] <http://www.ebay.com/>.
- [FF] <http://www.firefly.com/>.
- [FR93] D. Friedman and J. Rust, editors. *The double auction market: Institutions, theories and evidence*. Addison-Wesley, Reading, MA, 1993.
- [FSJ98] P. Faratin, C. Sierra, and N. R. Jennings. Negotiation decision functions for autonomous agents. *Int. Journal of Robotics and Autonomous Systems*, 24(3-4):159–182, 1998.
- [GMM98] R. Guttman, A. Moukas, and P. Maes. Agent-mediated electronic commerce: A survey. *Knowledge Engineering Review*, 13(2):147–159, 1998.
- [Gol92] R. Goldblatt. *Logics of Time and Computation, Second Edition, Revised and Expanded*, volume 7 of *CSLI Lecture Notes*. CSLI, Stanford, 1992. Distributed by University of Chicago Press.
- [HS69] J. Howard and J. Sheth. *The theory of buyer behavior*. Wiley, 1969.
- [JAN] <http://www.jango.com/>.

- [JPSF00] N. R. Jennings, S. Parsons, C. Sierra, and P. Faratin. Automated negotiation. In *Proceedings of the 5th International Conference on Practical Application of Intelligent Agents and Multi-Agent Systems (PAAM-2000)*, Manchester, UK, 2000.
- [KL89] K. Kuwabara and V. R. Lesser. Extended protocol for multi-stage negotiation. In *Proceedings of the Ninth Workshop on Distributed Artificial Intelligence*, Rosario, Washington, September 1989.
- [KM90] T. Kreifelts and F. von Martial. A negotiation framework for autonomous agents. In *Proceedings of the Second European Workshop on Modeling Autonomous Agents and Multi Agent Worlds*, Paris, France, August 1990.
- [KSE98] S. Kraus, K. Sycara, and A. Evenchik. Reaching agreements through argumentation: a logical model and implementation. *Artificial Intelligence*, 1-2(104):1–69, October 1998.
- [Lee96] L. Chi-Hang Lee. *Negotiation Strategies and their Effect in a Model of Multi-Agent Negotiation*. PhD thesis, Department of Computer Science, University of Essex, July 1996.
- [Nic66] F. Nicosia. *Consumer Decision processes: marketing and advertising implications*. Prentice Hall, 1966.
- [OFJ99] E. Oliveira, J. M. Fonseca, and N. Jennings. Learning to be competitive in the market. In *Proceedings of the AAAI Workshop on negotiation*, Orlando, USA, July 1999.
- [OFSG99] E. Oliveira, J. M. Fonseca, and A. S. Steiger-Garçon. Multi-criteria negotiation on multi-agent systems. In *Proceedings of the First International Workshop of Central and Eastern Europe on Multi-Agent Systems*, St. Petersburg, Russia, June 1999.
- [OR94] M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, Cambridge, Massachusetts, 1994.
- [OR00] E. Oliveira and A.-P. Rocha. Agents advanced features for negotiation in electronic commerce and virtual organisation formation process. In C. Sierra and F. Dignum, editors, *This book*. Springer Verlag, June 2000.

- [PC00] S. Paurobally and J. Cunningham. Formal models for negotiation using dynamic logic. In C. Sierra and F. Dignum, editors, *This book*. Springer Verlag, June 2000.
- [Per] <http://www.personalogic.com/>.
- [QXL] <http://www.qxl.com/>.
- [Res99] ActivMedia Research. Real numbers behind net profits 1999. Sixth annual study, New Hampshire, June 1999.
- [RNSP97] J. Rodriguez, P. Noriega, C. Sierra, and J. Padget. FM96.5 A Java-based Electronic Auction House. In *Proceedings of 2nd Conference on Practical Applications of Intelligent Agents and MultiAgent Technology (PAAM)*, pages 207–224, London, UK, April 1997.
- [RZ94] J. S. Rosenschein and G. Zlotkin. *Rules of Encounter: Designing Conventions for Automated Negotiation among Computers*. The MIT Press, Cambridge, MA, 1994.
- [San99] T. Sandholm. Distributed rational decision making. In G. Weiss, editor, *Multiagent Systems: A Modern Introduction to Distributed Artificial Intelligence*, pages 201–258. MIT Press, 1999.
- [SB98] R. S. Sutton and A. G. Barto. *Reinforcement Learning – An Introduction*. MIT Press, 1998.
- [SJNP98] Carles Sierra, Nick R. Jennings, Pablo Noriega, and Simon Parsons. A framework for argumentation-based negotiation. In M. P. Singh, A. Rao, and M. J. Wooldridge, editors, *Intelligent Agents IV (LNAI Volume 1365)*, pages 177–192. Springer-Verlag: Berlin, Germany, 1998.
- [TAT] <http://ecommerce.media.mit.edu/Tete-a-Tete/>.
- [Vic61] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.
- [VJ00] N. Vulkan and N. R. Jennings. Efficient mechanisms for the supply of services in multi-agent environments. *International Journal of Decision Support Systems*, 28(1-2):5–19, 2000.

- [VP00] M. Vetter and S. Pitsch. Towards a flexible trading process over the internet. In C. Sierra and F. Dignum, editors, *This book*. Springer Verlag, June 2000.
- [WJ95] M. Wooldridge and N. R. Jennings. Intelligent agents: theory and practice. *Knowledge Engineering Review*, 2(10):115–152, 1995.
- [WWW98] P. R. Wurman, M. P. Wellman, and W. E. Walsh. The Michigan Internet AuctionBot: A configurable auction server for human and software agents. In K. P. Sycara and M. J. Wooldridge, editors, *Proceedings of the 2nd International Conference on Autonomous Agents (Agents'98)*, pages 301–308, New York, May 9–13, 1998. ACM Press.