

# Efficient Mechanisms for the Supply of Services in Multi-Agent Environments

Nir Vulkan

Department of Economics

Bristol University

and the Centre for Economic

Learning and Social Evolution (ELSE)

University College London, London WC1E 6BT, UK

Email `N.Vulkan@ucl.ac.uk`

Nicholas R. Jennings

Department of Electronic Engineering

Queen Mary and Westfield College

University of London

Mile End Road

London E1 4NS

Email `N.R.Jennings@qmw.ac.uk`

## Abstract

Auctions provide an efficient way of resolving one-to-many negotiations. This is particularly true for automated agents where delays and long communications carry negative externalities. A properly designed auction, tailored to the specific needs of the relevant multi-agent system, can significantly improve its performance. In this paper we focus on the specific problem of service allocation amongst autonomous, automated agents, within the context of the ADEPT project which concerns the BT (British Telecom) business process of providing a quote for designing a network for a customer.

The main contributions of this paper are threefold: First, we show how an English auction can be modified for services, which are multi-dimensional private value objects. Second, we show how, under certain conditions, auctions can be arranged by the service providing agents, in the cases where the service seeking agents fail to do so. We consider the incentives of all participants, and show how such an arrangement can be in their best interest. Finally, by examining our results for what is, essentially, an application of game-theory and mechanism design to an existing application, we draw some general conclusions on how such concepts can be operationalized in automated agents.

## 1 Introduction

Autonomous agents are increasingly being required to operate in open, distributed systems comprising multiple problem solvers with competing objectives. In such situations, agents need to interact with one another in order to procure services and to manage their action interdependencies. As the agents are autonomous and because they often represent different organizations or individuals, the predominant form of interaction is negotiation. Thus, the agents decide for themselves which actions to pursue, for whom and under what terms and conditions. Only if they are convinced of the value of a particular course of action will they act toward its achievement.

Against this background, work on negotiation is clearly an essential component of multi-agent systems research and

development. There are, however, many different types of negotiation scenario which can be encountered (variables include the number of agents involved in the process, whether the negotiation is cooperative or competitive, and whether the agents repeatedly encounter one another) and each variation typically requires a different model. Moreover, there are a number of fundamentally different methods of approach to modelling negotiation — emanating from fields such as Artificial Intelligence (e.g. Kraus (1997), Mueller (1996)), Social Psychology (e.g. Pruitt (1981, 1993), Raiffa (1982)), and Economics (e.g. Kreps (1990), Binmore (1992)). What this melting pot of ideas shows clearly is that there is no universally applicable model and no universal best technique.

Here we are interested in a particular class of negotiations — namely, those where there is one buyer and many potential sellers. For this set of scenarios, we seek to devise a negotiation model which is fast and efficient and which can be easily implemented in a computational system. We chose to base our model on auction theory, since this is widely recognised by economists as the most efficient way of resolving one-to-many negotiations [see, for example, Moore (1992)]. This is particularly true for automated agents where delays and long communications can carry negative externalities, in addition to their usual direct negative effect on agents' utilities.

In this paper we focus on the specific problem of service allocation amongst autonomous, automated agents. The negotiation work is motivated by, and demonstrated on, a real-world problem from the domain of business process management. The particular problem we consider is taken from the ADEPT project (Jennings *et al* (1996)) and concerns an exemplar BT (British Telecom) business process for providing a quote for designing a network for a customer.<sup>1</sup>

In this application, autonomous agents represent the distinct actors involved in the process — thus, there are agents which represent the customer service division, the legal department, the design department and so on. These agents are conceptualized as providing services to one another. In order to be able to vend a particular service, an agent often requires an acquaintance to provide a related sub-service. Since the agents are autonomous, the means through which such sub-services are provisioned is negotiation. In this case,

---

<sup>1</sup> Although couched in terms of a particular business process management problem, we believe the scenario is common to a broad range of domains in which agents negotiate over the terms and conditions under which task allocation will occur (see Sierra *et al* (1997) for more details). Thus, the solution proposed in this paper is not just limited to the specific application we describe.

negotiation amounts to determining which agent should provide the desired sub-service and under what terms and conditions. One of the key elements of this application is a one-to-many, competitive negotiation for out-sourcing a particular part of the business process. Since this interaction is on one of the business process' main paths, it requires an efficient solution if the overall application is to be effective. Hence the interest in finding the most appropriate negotiation model.

The main contributions of this paper are threefold. First, we show how a standard English auction protocol can be modified for services, which are multi-dimensional private value objects. Second, we devise a novel protocol which, under certain conditions (which we specify), allows auctions to be arranged by the service providing agents, in the cases where the service seeking agents fail to do so. We consider the incentives of all participants, and show how such an arrangement can be in all their best interests. Finally, by examining our results for what is, essentially, an application of mechanism design to a computational setting, we draw some general conclusions on how the key concepts can be operationalized in automated agents.

The remainder of the paper is organized in the following way. Section 2 briefly describes the ADEPT system, and in particular focuses on the service allocation requirements dictated by our real-world application. Section 3 provides some background on auction design and outlines criteria for what is, and what is not, an efficient protocol. Section 4 describes the protocol for a multi-dimensional version of an English auction, designed by us for the ADEPT system. In this section we also provide some general results regarding the outcome and efficiency of our protocol. In section 5, a "pre-auction" agreement, organized by the service-providing agents, is described. We give an example of such a protocol and show when it could be used to increase efficiency. Section 6 discusses the application of our analysis to the design of interacting agents in general. Section 7 presents related work and section 8 concludes.

## 2 The ADEPT System

The ADEPT system provides a general framework and architecture for business process management which has been deployed in a range of real-world settings. Here we focus on one of these applications – namely, that of providing a quote for designing a network which provides particular services for a customer. The overall process receives a customer service request as its input and generates as its output a quote specifying how much it would cost to build a network to realize that service. The complete application involves up to six agent types (Jennings *et al* (1996)); however here we focus on the negotiation between only two types. The agents we consider are the *customer service division agent* which is responsible for managing the quote process for BT and the *vet customer agents* which are responsible for providing the out-sourced service of checking the customers who require quotes. The vetting of customers is an important aspect of the business process because if a customer is deemed unsatisfactory (for whatever reason) then a quote will not be provided.

There are a number of external organizations (typically between 5 and 10) who can provide the vet customer service to BT and they compete against one another to win the vet customer contract. In more detail, the negotiation between the customer service division agent (the *service-seeking agent*) and the vet customer agents (the *service-*

*providing agents*) exhibits the following properties:

(1) The vet customer service is provided by more than one agent. Some of the agents offer identical services, while others offer services which vary along several dimensions (*e.g.* quality, price, availability, etc.).

(2) The customer service and vet customer agents represent different organizations and hence the negotiation is competitive in nature.

(3) The set of vetting agents remains reasonably stable (hence negotiators repeatedly encounter one another), but new agents can enter the market and existing ones can leave.

(4) Negotiations range over a number of quantitative (*e.g.* price, duration, and cost) and qualitative (*e.g.* type of reporting policy, and nature of the contract) issues. Each successful negotiation requires a range of such issues to be resolved to the satisfaction of both parties. Agents may be required to make trade-offs between issues (*e.g.* faster completion time for lower quality) in order to come to an agreement.

(5) As the agents are autonomous, the factors which influence their negotiation stance and behaviour are private and not available to their opponents. Thus, agents do not know what utilities their opponents place on various outcomes, they do not know what reasoning models they employ, they do not know their opponent's constraints and they do not know whether an agreement is even possible at the outset (*i.e.* the participants may have non-intersecting ranges of acceptability).

(6) Time is an important consideration. Timings are important on two distinct levels: (i) the time it takes to reach an agreement must be reasonable; and (ii) the time by which the negotiated service must be executed is important in most cases and crucial in others. The former means that the agents should not become involved in unnecessarily complex and time consuming negotiations – the time spent negotiating should be reasonable with respect to the value of the service agreement. The latter means that the agents sometimes have hard deadlines by which agreements must be in place (this occurs mainly when multiple services need to be combined or closely coordinated).

Given this set of requirements, we need to design an appropriate negotiation protocol (mechanism) which ensures services are allocated efficiently and effectively. In our case, this involves taking the mechanism which best fits our needs, adapting it to give the desired properties and degree of efficiency, and then determining how it can be used by computational agents.

## 3 Designing Efficient Mechanisms

This section discusses the issues involved in designing an efficient auction protocol for automated agents operating in an open multi-agent system. This involves defining the criteria against which the protocol will be judged (section 3.1), determining those issues that arise over and above the standard ones because we are dealing with software agents not humans (section 3.2), and indicating how the uncertainty inherent in our negotiation scenarios can be dealt with (section 3.3). Finally, given this background, we motivate our choice of an English Auction Protocol to handle the ADEPT scenario (section 3.4).

### 3.1 Evaluation Criteria

In order to compare two negotiation mechanisms, some assumptions about the behaviour of agents are necessary. In

the economic literature, the usual method of comparing different negotiation mechanisms (or protocols) is to examine some facet of the set of Nash-Equilibria of the games specified by these mechanisms (Moore (1992)). A Nash-Equilibrium is a state where each of the players best-respond to the collection of the other players' strategies. For example, in a protocol where two negotiating agents are required to simultaneously demand what share they should be getting from a given sum of money, a pair of demands constitutes a Nash Equilibrium if and only if, the shares demanded sum up to one (because the best response to any suggested share, is to demand one minus that share). Implicitly, getting to a Nash equilibrium generally requires that players are rational, in the sense that they act in such a way to maximize their goals, and that this rationality is common knowledge (for a more detailed discussion see Binmore (1992), or Moore (1992)). That is, each agent knows that its opponent is a utility maximizer, that its opponent knows that it is a utility maximizer, that it knows that its opponent knows, ad infinitum. However, as we explain later in this paper, some mechanisms, like the ones suggested by us, make more moderate assumptions on the rationality of agents which are better suited to computationally bounded systems.

When considering notions like utility maximisation it is important to distinguish between *private* and *common value* objects. The worth of a private value object to a bidder depends only on his own preferences. In contrast, the value, to a bidder, of a common value object depends on the preferences of the other bidders<sup>2</sup>. A service in a multi-agent environment has private value if the only utility it carries for the service-seeking agent is from consuming it. Conversely, if the agent is allowed to re-sell the service, then it has common value. Since information regarding the credit rating of a potential customer cannot (legally) be distributed, ADEPT's customer service agent cannot re-sell the service. In this paper, we therefore restrict attention to private value auctions. However, note that the outcomes of private and common value auctions differ in general, and some of the results in this paper will need to be modified for a multi-agent system where re-sale of services is allowed.

With these basic constructs in place, we can now return to the issue of mechanism evaluation. For our purposes, there are two important criteria that need to be considered. Firstly, the expected revenue, or utility, of the auctioneer. In our setting, this equates to the utility the buyer receives, in equilibrium, from the service specified by the winning bid. The second metric relates to the mechanism's allocative efficiency, or efficiency in general. In our setting, this translates to comparing the sum of the utilities of all bidders, or to the question of whether the winner is the service-provider which could offer the best service for the buyer whilst still profiting from such an agreement. If communications are expensive (we can think of the direct costs of communicating between agents, or a slight preference for agreements closer in time), efficient also means fast (see Binmore (1992)).

## 3.2 Multi-Agent Specific Considerations

There are several ways in which mechanism design for multi-agent systems differs from that for multi-person environments (see, for example, Rosenschein and Zlotkin (1994), Sandholm and Lesser (1995), Sandholm (1996), and Binmore & Vulkan (1997)). In our case the following differences are of relevance. On the positive side, the environment

<sup>2</sup>A third possibility is when these values depend on both private and common preferences (known as *correlated value* auctions).

is controlled and we do not have to worry about signaling problems as in the case of human players. Consider, for example, the recent U.S. bandwidth auctions (see McMillan (1994) for more details about these auctions). Here companies compete for licenses in different cities and states. A crucial feature of the design of these auctions is that bidders cannot communicate to each other the destinations they would like to have. It was reported that in one of the recent auctions, one of the bidders submitted a surprisingly low bid (low enough to ensure that it will not win), and that the last few digits of amount bid were identical to the telephone code of the area it was after. Although this was not coordinated between bidders in advance, the signal was correctly interpreted by the other players, and the revenue generated was significantly lower than what was originally anticipated. At least with existing technology, computerized agents are unlikely either to be capable of generating such subtle signals or of interpreting them even if they are produced. Another positive difference is that mechanisms will be executed quickly. Moreover, as the number of agents is relatively small, we do not have to worry about timing of messages (cf large scale auctions over a distributed system, like the Internet (Rodriguez *et al* (1997))). On the negative side however, automated agents' knowledge bases are inherently bounded: they cannot initiate a mechanism unless they are specifically designed to do so. Also, the idea that agents always best-respond, which in many simple situations is reasonable for humans, is questionable for automated agents.

The other major difference when designing for a multi-agent system is that the system is open in nature. Traditional mechanism design analysis only compares the outcomes of a given set of mechanisms. In other words, mechanisms are taken as given. However, in an open system, where the type of future participating agents is still unknown, we may wish to consider the incentives of agents to *initiate* a mechanism. Specifically we consider (i) the conditions under which a service seeking or service provider agent will choose to initiate a particular mechanism, and (ii) the knowledge base automated agents will be required to have in order to decide that a certain mechanism is desirable in a given situation. We address the former question in sections 4 and 5, and the latter in section 6.

## 3.3 Coping with Uncertainty

In most multi-agent systems, service-seeking and service-providing agents are uncertain about the valuation systems of their opponents.<sup>3</sup> In such situations the behaviour of each agent is, in general, determined by its beliefs about the behaviour of the other agents. More specifically, in our setting, the behaviour of a bidding agent clearly depends on its beliefs regarding the valuation systems of other bidders. An immediate consequence is that the outcome of the chosen mechanism/protocol may depend on the collection of these beliefs: outcomes are a function of beliefs. This has to be taken into account when we evaluate the efficiency (or any of the other criteria mentioned above) of such a mechanism. Several evaluation premises could be considered: given the distribution of beliefs (or the modeller *a priori* beliefs about beliefs) we could compute an average outcome, with respect to beliefs, or we could restrict attention to worst-case

<sup>3</sup>In fact, if there was no uncertainty, the whole negotiation process could efficiently be reduced to a simple choice function, with participants having no incentive to deviate from what is dictated by that choice function.

scenarios<sup>4</sup>. Since we need more assumptions to discuss averages, we focus in this paper on the latter. However, some of the results, as we point out in section 4, will still go through in the average-outcome case. By worst-case scenario, we have in mind something similar to the notion of worst-case algorithmic complexity. Thus we will consider the outcome when participants' beliefs are such that they minimize the efficiency of the mechanism we will put forward.

However, in many situations it is possible to construct mechanisms where the outcome does not depend on beliefs. A *belief-independent mechanism* exists when outcomes depend on the behaviour of a sub-set of agents (at least one, and possibly all) which have a *dominant strategy*. A dominant strategy is one which yields a (expected) payoff which is higher than the other strategies *whatever* the behaviour of the other players and the state of the world. For example, in the prisoners' dilemma a player prefers to defect if the other player cooperates, and prefers to defect if the other player defects. Defecting is, therefore, a dominant strategy.

Dominant strategies have several desirable properties when applied to automated agents. First, they permit a much weaker form of rationality — a player should never play a strategy which is dominated. This last statement can be easily represented in the knowledge base of an automated agent. In fact, assuming that this simple fact is commonly known to all agents ensures that agents will play the equilibrium in a large set of games (Pearce (1984)). If agents can reach the equilibrium by elimination of dominated strategies, then this can be done with relatively little counter-speculation about what the other players will do. Given the complexity of knowledge representation induced by such counter-speculations, mechanisms which rely on dominant strategies are clearly desirable. Second, think of the behaviour of rational agents (*i.e.* utility maximizers) if, for some reason, the system is invaded by one or more “irrational” agents. A mechanism which relies only on dominant strategies will be robust to such an event, since optimal actions are independent of the behaviour of the other players. However, this is not the case in situations where what is optimal depends on the actions of others. In the latter case, agents will hold and update beliefs regarding the rationality of the other agents, and will act accordingly (by, for example, taking advantage of the irrational agents). This kind of behaviour in turn, could reduce the overall efficiency of the system. For these reasons, we seek to design an auction protocol for the ADEPT scenario based upon agents selecting dominant strategies.

### 3.4 Choosing a Particular Auction Protocol

Perhaps the most common auctions mechanisms are (i) a Dutch auction, where the auctioneer continuously lowers the price until a bidder accepts the price and the auction terminates, (ii) an English (or First-Price, Open Cry) auction, where the auctioneer raises the price until there is only one taker for that price, (iii) a first-price sealed bid auction, where bids are made simultaneously and the highest bid wins the auction and where that bid is taken as the agreed price, and (iv) a Vickrey auction, where simultaneous sealed bids are made. The winner is the agent with the highest bid, but he only pays the second highest bid. Although it is possible that, under some conditions, all four auction mechanisms are revenue equivalent for a private-value object (e.g. Moore (1992)), only two (English and Vickrey auctions) rely

<sup>4</sup>Alternatively, any combination of the two measurements can be considered.

on agents playing their dominant strategies. As indicated in section 3.3, this has important conceptual and pragmatic advantages.

The reason for choosing the English auction over the Vickrey auction is that it is more robust with respect to changing circumstances. Suppose now, that in some situations, services *can* be re-sold (or more generally, that there are some externalities (e.g. the fact that agent  $i$  had won the auction, has an indirect effect on agent  $j$ 's utility) in the process of service provision). The object (*i.e.* the service) is no longer of private value, and the equivalence theorem no longer holds. In particular, the choice of a Vickrey auction is no longer likely to be optimal (Moore (1992)). But an English auction is still revenue superior in a very large class of situations, because the information revealed during the auction ensures that bidders will move closer to their reservation prices (Moore (1992)).

Given this background, in the next section we describe a mechanism for service allocation suitable for the ADEPT application, based on a multi-dimensional English auction. We show that our choice of mechanism relies on dominant strategies, promotes truth revealing from the side of the service-seeking agent, and can be executed quickly.

## 4 A First-Price, Open-Cry Auction Protocol for ADEPT

In this section we design a new protocol for service allocation in the ADEPT scenario. Our protocol represents an extension of the standard English auction in that it is multi-dimensional. We provide a specification of the implementation of such a protocol, and re-derive the usual efficiency results. In particular, we show that our protocol will choose the service provider which can make the best offer, in terms of the buyer's utility. We also show that the buyer will then choose to report its true utility (*i.e.* the way it weights the different attributes of the service). Finally, we show that the protocol maximises the worst-case outcome for the buyer (see section 3.3).

Multi-dimensional auctions add new complexity to the auctioning process, because sellers (and buyers for that matter) now typically have different preference weightings. This implies that the service provider which can make the best offer is not necessarily the one that offers the lowest price. Rather it is the one that offers the combination most suitable for the buyer's needs. That is, for fixed prices, different buyers will prefer different service providers. This is in sharp contrast to the one dimensional case, where the best-price is preferred by all. Still, as we demonstrate (section 4.2), it is possible to get similar efficiency results for a suitably modified version of the English auction protocol. We, therefore, begin with definitions of multi-dimensional services and preferences (section 4.1).

### 4.1 Defining the Protocol

**Services:** We describe a service as a pair:  $(p, \bar{s})$  where  $p$  is a real number corresponding to the price of the service, and  $\bar{s}$  a vector of real numbers  $\bar{s} = (s_1, s_2, \dots, s_N)$  where  $s_1 = q$  is a measurement of the quality of the service<sup>5</sup>, and  $s_2$  to  $s_N$  are optional additional dimensions of the same service (such as time before service can be provided, level of support, etc.).

**Preferences of the service-seeking (SS) agent (or “Buyer”):** We assume that the ss-agent's preferences can

<sup>5</sup>We assume that the quality of a service can be described by a real number and that this measurement is commonly known to all players.

be described by a *quasi-linear utility function*<sup>6</sup>  $u_B(\bar{s}) - p$ : A continuous and differentiable function which, we assume, is increasing with the quality of the service agreed,  $\frac{du_B}{ds_1} > 0$ . The ss-agent's spending on the service is restricted by some real number  $p_0$ . It makes economic sense, certainly within the context of ADEPT, to assume that  $u_B$  is concave. In consumer theory, this is known as the law of diminishing marginal rates of substitution: to hold utility constant, diminishing quantities of one attribute of the service must be sacrificed to obtain successive equal increases in the quantities of the other attributes (see, for example, Binmore (1992) or any microeconomic textbook for more details).

**Cost structure of service providers (SP) (or "Sellers" or "Bidders"):** The costs to a sp-agent  $i$ , from a service  $\bar{s}$  and a price  $p$  can be expressed by a continuous and differentiable function  $c_i(\bar{s}) + p$ , with  $\frac{dc_i}{ds_1} < 0$ . Moreover, we assume that the cost functions faced by service providers are such that, for each  $s_j$  where  $\frac{du_B}{ds_j} < 0$  then  $\frac{dc_i}{ds_j} \geq 0$  and vice versa for the service sellers. That is, the sp-agent's preferences move in opposite directions to those of the buyer. This assumption is not restrictive: suppose that there are some variables that buyers and sellers prefer more (alt. less) of. Then, those variables can be excluded from the analysis above, with the outcome remaining unchanged. We assume further that  $u_i(\cdot)$  are convex (or diminishing marginal returns of the cost function). This assumption is taken for the same reasons, and in the same vein, as the previous assumption on diminishing rates of substitution.

#### The Protocol:

*I. Initiation:* The process begins with an announcement from the ss-agent. An announcement consists of the following list  $\{FPOC, \hat{u}_B(\bar{s}), \hat{p}_0, mi, T, X\}$  where *FPOC* describes the type of auction which the ss-agent is about to use, in our case, FPOC or first-price open-cry;  $\hat{u}_B$  is the buyer's declared utility function (but in general  $\hat{u}_B$  need not be the real utility  $u_B(\bar{s})$ );  $\hat{p}_0$  is the maximal amount the buyer is willing to pay for the service (and again,  $\hat{p}_0$  can be in general different from  $p_0$ ); *mi* (optional) is the minimal acceptance level for offers, i.e. offers  $\bar{s}_i$  such that  $\hat{u}_B(\bar{s}_i) < mi$  will not be considered;  $T$  is the maximal time (in seconds) the buyer will wait for a new offer before accepting the existing one; and finally,  $X$  is the percentage figure by which a new offer has to exceed the last offer in order to be considered. Formally, given the last acceptable offer  $\bar{s}_t$ , a new offer,  $\bar{s}_{t+1}$  will be considered iff  $\frac{\hat{u}_B(\bar{s}_{t+1})}{\hat{u}_B(\bar{s}_t)} > 1 + X$ . Note that although the ss-agent is allowed to lie about the values of  $u_B(\bar{s})$  and  $p_0$  we will show that this will not happen in equilibrium.

*II. Auction:* Once sp-agents receive the announcement they can submit offers — where an offer is simply a vector  $\bar{s}$  as defined above. An offer is accepted providing that (1) the protocol did not terminate, and (2) the offer exceeds the last offer by at least  $X$  percent. When an acceptable offer is submitted, it is made public by the buyer (who acts as the auctioneer in the more conventional version of an English auction), together with the identity of the bidder.

*III. Termination:* The auction terminates  $T$  seconds after the last acceptable offer was made.

#### End of Protocol.

<sup>6</sup>It is assumed that these preferences are identical to those of the agent's user, or that the user never lies to his own agent. The idea behind this assumption is that the agent will do the lying on behalf of its user, if this could be beneficial. See Binmore and Vulkan (1997) for a discussion of this assumption.

The role of  $T$  should be obvious from the above termination condition. Of course,  $T$  has to be long enough for an interested sp-agent to be able to carry out its calculations and submit an offer. Since this is the only reason why it has to be long, we set  $T$  to be the smallest time interval which satisfies the above condition<sup>7</sup>.

$X$ , like  $T$ , is derived from practical considerations: Typically, the ss-agent has to get the service before some deadline (section 2). Although the agent does not discount the future by some continuous factor (as in most one-to-one bargaining models, where time discounting enters directly to the agents' utility functions), nevertheless it prefers mechanisms where agreements can be reached in a known, finite time. If bidders are allowed to increase existing offers by any amount, a situation where offers increase by a decreasing percentage cannot be ruled out. In theory, these increments can become smaller and smaller and convergence might be extremely slow (although it is not in the interest of any of the players). Fixing  $X$  will clearly prevent such undesirable scenarios. Note, though, that fixing  $X > 0$  may mean that the ss-agent could lose out some of its surplus in the following scenario: Suppose that the current offer is  $\bar{s}^T$ . Suppose further, that some bidder,  $i$ , is able to offer a slightly better service to the buyer,  $\bar{s}^i$ , where  $u_B(\bar{s}^i) > u_B(\bar{s}^T)$ , but  $\frac{u(\bar{s}^i)}{u(\bar{s}^T)} < 1 + X$ . Then  $i$  will not bid and, if  $\bar{s}^T$  is accepted, the buyer could have been better-off if  $i$  bid.

## 4.2 Analysis of the Protocol

The following simple Lemma is useful for the rest of our analysis:

**Lemma 1** Under the above assumptions, there exists a unique solution,  $\bar{s}^i$ , to the following constraint maximization problem for sp-agent  $i$ :  $\max u_B(\cdot)$  subject to  $c_i(\cdot)$ , and  $p \leq p_0$ .

**Proof.** Let  $\mathcal{S}_{\min} = \{\bar{s} \mid c_i(\bar{s}) = -p_0\}$ . For the given price  $p_0$ ,  $\mathcal{S}_{\min}$  represents the set of all services for which service provider  $i$  is indifferent between providing and not-providing for that price. From continuity,  $\bar{s}^i$  must lie on  $\mathcal{S}_{\min}$  (otherwise we should be able to move closer to  $\mathcal{S}_{\min}$  while  $i$  still makes positive profits). For every real number  $k$ , consider the set  $\{u_B(\bar{s}) = k\}$ , what is known in economics as the buyer's indifferent curve, or in our setting the set of all services that are worth exactly  $p + k$  to the buyer. Given the convexity/concavity assumptions, there exists a unique real number  $\bar{k}$  such that  $\mathcal{S}_{\min}$  is tangential to the indifferent curve  $u_B(\bar{s}) = \bar{k}$ . See figure 1 for an illustration of these curves in the case where  $\bar{s}$  is two dimensional. Given the direction of the buyer's preferences the point where the two above curves are tangential maximizes his utility. This unique point on the services' space is therefore  $\bar{s}^i$ .  $\square$

Lemma 1 describes the best possible agreement, *to the buyer*, that service provider  $i$  could offer for the given price. However, to expect to extract such an agreement in equilibrium is unrealistic, as the provider is primarily concerned with minimizing his own costs (and not with maximizing the buyer's utility as in the proof of lemma 1), and will therefore try to win the auction with the service which does just

<sup>7</sup>Unlike a one-to-one bargaining situation, where delays plays a significant strategic role: by not accepting offers immediately, a player is signaling that he is patient, which in turn, increases his bargaining power. However, in our setting, if a provider wishes to offer more than the existing offer, it has no incentive to wait. See, for example, Admati and Perri (1987).




Figure 1: Buyers best possible agreement.

that. As we show next, the best that the buyer can do, is to extract from the winner, a service which is (slightly higher than) the second-best. Formally,

**Proposition 1** The auction will be won by sp-agent  $j$ , where  $u_B(\bar{s}^j) > u_B(\bar{s}^i) \forall i \neq j^8$ , and the winning bid,  $\bar{s}^*$  is such that  $u_B(\bar{s}^*) = (1 + X) \cdot u_B(\bar{s}^k)$ , where  $\bar{s}^k$  satisfies  $u_B(\bar{s}^k) > u_B(\bar{s}^i) \forall i \neq j, k^9$ .

**Proof.** (By reductio ad absurdum) Suppose, on the contrary, that agent  $l \neq j$  wins the auction. Since  $u_B(\bar{s}^j) > u_B(\bar{s}^i) \forall i \neq j$  by assumption, then  $u_B(\bar{s}^j) > u_B(\bar{s}^l)$ . From the definitions of  $\bar{s}^j$  and  $\bar{s}^l$ , sp-agent  $j$  receives a higher utility from winning with  $\bar{s}^j$  than from losing the auction, and  $l$  cannot submit a better bid than  $\bar{s}^l$ .  $j$  will therefore profitably deviate by bidding  $\bar{s}^j$  which will result in him winning the auction.

It should be clear that  $j$  has no incentive to bid anything which yields the buyer more utility than  $\bar{s}^*$  (recall our assumption that the buyer's preferences are opposed to those of the sp-agents. Any agreement which improves the utility of the buyer will therefore decrease the utility of the seller). Suppose now that  $j$  wins the auction with a lower bid,  $\bar{s}^{**}$  (where  $u_B(\bar{s}^{**}) < u_B(\bar{s}^*)$ ). Then sp-agent  $k$  could win (or, at a minimum, tie) the auction by bidding  $\bar{s}^k$ . Hence our result.  $\square$

In fact, in an English auction it is a dominant strategy for bidders to increase their bids (with respect to  $u_B$  until  $\bar{s}^j$ ). Proposition 1 shows, therefore, that the buyer can expect to extract the “second best” service from the bidders,

<sup>8</sup>Providing that  $u_B(\bar{s}^j) > \min$ . That is, providing that the “best” offer is acceptable. Otherwise, none of the offers is acceptable and the service will not be provided.

<sup>9</sup>That is, sp-agent  $k$  would have won the auction if  $j$  did not participate.

regardless of their beliefs. Given that, we show next that the ss-agent will prefer to report to the sp-agents his true utility and price constraint (whereas in one-to-one negotiations such a revelation is very unlikely, as revealing one's private information weakens its bargaining power).

**Proposition 2** In equilibrium, the ss-agent will announce  $\hat{u}_B(\cdot) = u_B(\cdot)$  and  $\hat{p}_0 = p_0$ .

**Proof.** Given the behaviour, in equilibrium, of the sp-agents, the ss-agent will receive utility  $\hat{u}_B(\bar{s}^*)$  which is clearly maximized by choosing  $\hat{u}_B(\cdot) = u_B(\cdot)$ . Suppose now that the service is agreed for  $p_1 < p_0$ . Then given the opposite direction of the cost functions and  $u_B$ , and the continuity assumption on both, there is an  $\epsilon$  environment which contains services with costs  $p$ ,  $p_1 < p < p_0$  for which  $u_B$  is strictly increased. The ss-agent is therefore better-off announcing  $\hat{p}_0 = p_0$ .  $\square$

Given the results from Propositions 1 and 2, we are now able to prove that the protocol is, in some sense, optimal:

**Proposition 3** The above protocol maximizes the beliefs-worst-case revenue for the buyer.

**Proof.** In the game where  $u_B$  and the  $c_i$  are commonly known, the equilibrium outcome cannot exceed  $u_B(\bar{s}^*)$ . Consider now any protocol  $P'(\bar{b})$  in which the outcome (may) depend on the collection of beliefs. Let  $P'(\bar{b}^T)$  denote the outcome of the protocol when all players place probability 1 on the true values of  $c_i$ . Clearly  $u_B(P'(\bar{b}^T)) \leq u_B(\bar{s}^*)$  because, given their beliefs, none of the bidders will have an incentive to offer a “better” service for a price  $p \leq p_0$ . By definition, belief-worst-case outcome,  $u_B(P'(\bar{b}^{w.c.}))$  is not greater than  $u_B(P'(\bar{b}^T))$ . Hence our result.  $\square$

In fact, with more work, it is possible to prove a stronger version of Proposition 3, namely that under a large class of distributions of beliefs (including the uniform distribution), the above protocol also maximizes the average-case outcome. Note, however, that in order to prove such claims, we need to make further assumptions about the risk-attitudes of sp-agents (because the utility of an average is equal, smaller or greater than, the average utility depending on the risk attitude of the agent). Still, it is easy to see from the proof of Proposition 3, that it will still work as long as the average belief does not imply over-optimism, in the sense that sellers are likely, on average, to under estimate the cost structures of their opponents.

To summarize, Propositions 1-3 show that it is in the best interest of the buyer to use the suggested protocol in order to ensure the best possible outcome for itself. No other mechanism, whether it is based on an auction or on direct negotiations, will yield a better outcome for the service seeker. Moreover, the mechanism chosen is robust against changing circumstances, and since it is expressed in multi-dimensional terms, is applicable for a large set of services.

## 5 A Pre-Auction Agreement Protocol

The previous section dealt with the case of an auction initiated by a service-seeking agent. Since the protocol provides an efficient method for the ss-agent to agree a service, it is likely that it will opt for an auction as it is typically better

off in an auction than in a situation of simultaneous negotiations with sellers.<sup>10</sup> Our goal in this section, however, is to show that if for some reason the service-seeking agent fails to initiate such a protocol, it might still be in the interest of the sp-agents to do so. Here we assume that the default interaction mechanism is direct negotiations. This is the case, not only in ADEPT, but in any other multi-agent system which allows for both one-to-one and one-to-many bargaining. We identify the conditions under which the sp-agents might decide that some sort of auctioning is required, and we discuss one particular mechanism which these agents can employ, and which does not hurt the buyers.

Consider the current state of affairs in ADEPT where the ss-agents are already endowed with a number of bargaining techniques. These agents are also familiar with the set of possible vet customer agents competing to provide them services. Quite often, the service which is agreed on does not differ much from the one that would have been agreed in the English auction described in the previous section. Moreover, in ADEPT a single service constitutes only a small fraction of the vet customer agents' turnover. The benefits to an sp-agent from not having to participate in an auction (namely the probability the agent attaches to the event of it winning the auction although it cannot offer the best possible service for that price, or the probability attached to winning the auction with a service which is less costly), are therefore small. Moreover, one-to-one negotiations are, almost by definition, longer and require more communications between the negotiating agents (in ADEPT, like in most multi-agent systems, negotiations are primarily a process of offers and counter-offers). The number of communication channels is limited, and long negotiations with one ss-agent could delay negotiations with other such agents, who might turn elsewhere in the meantime. From this, together with the additional assumptions on time discounting and deadlines, it should be clear that, for a fixed outcome, a fast mechanism is preferred. If, like in ADEPT, the benefit from direct negotiations is small, then an auction could be preferred by the sp-agents for even relatively small communication costs and time discounting<sup>11</sup>.

## 5.1 Defining the Protocol

In ADEPT, as in many other existing multi-agent systems, auctions are always initiated by the ss-agent. If an ss-agent, however good in negotiating deals on behalf of its user, does not know how to initiate an English auction, or if it fails to detect that it is appropriate, then costly one-to-one negotiations will continue. Denote this as Condition 1. Next, denote by Condition 2 the situation where the fine details of the service need to be finalized *after* one sp-agent is selected (this is typically the case where the setting of the final details of the bid are expensive). For situations where either of the above conditions hold, the following mechanism could prove beneficial:

### The “pre-auction” protocol:

<sup>10</sup>Binmore (1985) developed a model of simultaneous one-to-one bargaining which shows that, under certain conditions, the same outcome as an auction can be attained. However, this procedure is slower than an auction and therefore less efficient

<sup>11</sup>The exact calculation depends on the probabilities agents attach to the different outcomes. If agents update beliefs according to experience, then those beliefs will converge to the actual statistics (or the “truth”). If ss-agents do get “good deals” for themselves in direct negotiations, sp-agents will, from some point, prefer an auction.

*I. Initiation:* An sp-agent approached by an ss-agent (where either conditions (1) or (2) apply), announces a “pre-auction” protocol to the other sp-agents.

*II. Pre-Auction:* sp-agents compete using an English auction (as in section 4), with the difference that the maximal price is taken as  $(1 - A) \cdot p_0$ , where  $A$  is small (typically 0.01). Denote by  $\bar{s}^*$  the winning bid, and by  $i$  the index of the winning sp-agent.

*III. Negotiations:*  $i$  makes a take-it-or-leave-it offer to the ss-agent of service  $\bar{s}^*$  for a price of  $p_0$  (in equilibrium the offer will be accepted). All other sp-agent do not negotiate with the ss-agent.

*IV. Insurance:*  $i$  then pays the other  $n$  sp-agents,  $A \cdot p_0/n$  for their cooperation.

**End of Protocol.**

## 5.2 Analysis of the Protocol

For the above protocol we are able to prove the following results:

**Proposition 4** Given the above protocol; (i) sp-agents will participate in the pre-auction, (ii)  $i$  will make a take-it-or-leave-it offer  $\bar{s}^*$ , at a price  $p_0$  (iii) the ss-agent will accept, and (iv) all sp-agents, other than  $i$ , will not negotiate with the ss-agent.

**Proof.** Given that the sp-agent who is able to offer the best possible service (as in Lemma 1), wins the auction, none of the losing parties will have an incentive to start negotiations with the ss-agent. Given that the winning bid is  $\bar{s}^*$  (as defined in Proposition 1), the ss-agent's best response is to accept the take-it-or-leave-it offer immediately. The additional incentives (the insurance payment) are necessary to ensure that each sp-agent is better-off participating in the pre-auction. Finally, the proof regarding the winning sp-agent and the winning bid is identical to that of Proposition 1.  $\square$

In Proposition 4 we showed that in equilibrium, the pre-auction protocol will result in an efficient outcome. Next we show that under conditions 1 or 2, this provides sufficient incentives for the sp-agents to initiate such a protocol.

**Proposition 5** If communication is costly, or if sp-agents discount time, then the allocative efficiency will be increased by using the pre-auction mechanism.

**Proof.** Compare the outcome of the pre-auction protocol with that of direct negotiations (which may depend on the beliefs of bidders). From the proof of Proposition 1 it is clear that the allocative efficiency of bidders (excluding the ss-agent) cannot decrease by using a protocol with an efficient outcome. The utility of the ss-agent from direct negotiations is at most  $A \cdot u_B(\bar{s}^*)$  higher than it receives when the pre-auction protocol is used. But if either conditions 1 or 2 hold, then the utility of at least one (and possibly all) sp-agents is increased by using the protocol. Hence our result.  $\square$

In this section we showed how it is possible to overcome inefficiencies caused either by costly bidding, or by the failure of the ss-agent to initiate an efficient protocol. In such circumstances, it may very well become in the best interest of the service providers to restore efficiency. We show that

pre-auction mechanisms provide us with the means to do that. Moreover, our choice of mechanism guarantees also that the interests of the service seeking agent, in addition to the interests of the service providers, are being kept. This is important, because otherwise service seeking agents may oppose to the outcome of such a mechanism and will pursue better agreements independently, thus reducing efficiency even further. Note that the service seeking agent is slightly worse-off than in the protocol described in the previous section (because of the proportion  $A$  which is used to pay for the cooperation of the sp-agents). However, if  $A$  is small compared to the inefficiencies caused by not using a direct auction, then everyone may be strictly better-off by using our suggested pre-auction mechanism. Because automated agents are inherently bounded, either conditions 1 or 2, as described above, are likely to occur in open, multi-agent environments. The pre-auction protocol described here could, therefore, provide a simple and efficient way to resolve such problems. The ideas presented in this section represent an important divergence from traditional mechanism design, where it is assumed that once the protocol has been set all agents behave optimally.

## 6 Insights for the Design of Automated Agents

This section highlights the general issues involved in taking the theoretical results from mechanism design studies and using them to structure the design and implementation of practical agents. While these insights obviously apply to the ADEPT application discussed herein, they also have a much wider scope. In fact, we contend the following thoughts are relevant for all types of negotiation situation where the designer adopts a game theoretic start-point.

*Responding to existing protocols:* One of the main contributions of game-theory is that it provides a classification of negotiation situations based on their underlying strategic form (see, for example, Kreps (1990)). This fact can be exploited by the designers of bounded negotiating agents. Rather than building an elaborate reasoning mechanism which attempts to best-respond to any unforeseen protocol (undoubtedly a doomed venture!), the designer can use the broad classification of protocols, based on game-theoretic ideas, to indicate the most efficient way of responding in a given situation. Thus, responding becomes a case-base retrieval process, rather than a reasoning from first principles idea. As our work suggests, negotiating agents will have to be familiar, at a minimum, with the main types of auctions. Moreover, to successfully participate in auctions, bidding agents need to distinguish between private and common value objects. However, given the equivalence between the outcomes of many auction types (see Binmore (1992)), the number of protocols agents will need to know, might be comparatively small (we estimate in the order of 10). The database of cases will include the name of the protocol, the best-response (which, in general, may depend on real-time input, such that can be deduced from the bidding of other agents), and the anticipated outcome of the agent from following his best response in the given protocol. Finally, following what we know about the limitation of auction protocols (see, for example Sandholm (1996)), we suggest that the entry will also include a list (possibly empty) of conditions which violate the applicability of the best-response.

*Initiating a protocol:* Consider the expectations of a negotiating agent regarding the service agreed at the end of the negotiation process. Using the data-base suggested above,

the agent can check, at regular intervals, whether he expects to be better-off by using any of the mechanisms from that list. If the answer is negative, then clearly the agent has no incentive to change the negotiation protocol. But what if the answer is positive? Here the agent needs to know whether it is, technically, possible for it to initiate the preferred protocol. As we showed in section 5, the service-seeking agent is often better-off auctioning its custom. We also showed that, in certain circumstances, the service-providing agents will prefer an auction to one-to-one negotiations. However, a service providing agent cannot force the service-seeking agent to initiate the English auction protocol. But, in such situations, a different protocol, the “pre-auction” protocol, is technically feasible, and will have a very similar desirable outcome. We believe that it is possible for our other results to be further generalized in a similar fashion.

*Choosing protocols for multi-agent systems:* From the point of view of the designer of a multi-agent system who is concerned primarily with the efficiency of negotiations, we identify two important selection criteria. First, the choice of protocols should be robust with respect to the degree of rationality of agents. On the extreme, we should prefer protocols which are robust against the presence of non-maximizer agents. But as we explained in section 3, even if agents do maximize the utility functions of their users, bounds on their knowledge representation capabilities could prevent them from playing the Nash equilibrium of an unforeseen game. Fortunately, protocols which rely solely on dominant strategies (or on iterated elimination of dominated strategies) require a much simpler knowledge base, and are robust against the presence of bounded and irrational agents. At the risk of repeating ourselves, we again stress that it is therefore *extremely* important that such protocols will be used in multi-agent systems. Second, protocols which are robust with respect to changes in circumstances should be preferred. In this paper we chose an English auction over to its equivalent Vickrey auction because the latter was shown not to be optimal when the service is no longer of private value. Although this is not the case in the existing ADEPT framework, this choice allows us more flexibility, arguably the most important feature of a successful multi-agent environment.

## 7 Related Work

Rosenschein and Zlotkin (1994) were the pioneers of using Game Theory as a framework for modelling multi-agent interactions. They also promoted the use of mechanism design theory for specifying multi-agent interactions. Our first protocol can, therefore, be seen as continuing from their line of work. However our work takes mechanism design a stage further. Rather than using idealised, toy scenarios we have used it in a real-world context with all the concomitant constraints and difficulties that this brings. In a much simpler context, Varian (1996) is an excellent example of how second-price auctions can be generalized for a multi-purpose multi-agent system. In the same article, he points out the importance of mechanisms which rely on dominant strategies. As we indicated in section 3.3, this is particularly important for automated agents, where strategic-knowledge representation might otherwise be a problem. Moreover, sections 5 and 6 suggest that one cannot take the mechanism design literature at face value. Varian, like Rosenschein and Zlotkin, assumes that agents respond optimally to the given mechanism. We show that this assumption is problematic in the context of an open-for-all, multi-agent environment.



More generally, the economic literature on mechanism design and on auctions is vast and beyond the scope of this section. For a good survey of mechanism design see Moore (1992), or the relevant chapters of Binmore (1992). Varian (1996) and the references therein, provide a survey of second-price (or generalized Vickrey) auctions.

As we show in this paper, in order to adapt an existing mechanism to a given multi-agent system, it is necessary to try and envisage what problems will arise from the implementation. (Lee (1996) is an example for such considerations in the context of group negotiations). In particular, the system must be robust against non-optimizing agents, and should preferably rely on mechanisms where agents use dominant strategies. Specifically, we use a multi-dimensional version of an English auction for the supply of services in ADEPT. When auctioning a multi-dimensional object (or a multi-attribute service), knowledge of the reservation prices for the different attributes is not sufficient to determine a winner. One has to consider additionally, how the sp-agents weight the different attributes. This is somewhat similar to the idea of auctioning multi-units of a given good, a topic which received some attention in the economic literature (see, for example Moore (1992), Varian (1996), or Binmore (1992)). Here, each agent has a reservation price for each unit of consumption (e.g. agent  $i$  is willing to pay up to 5 dollars for the first unit of the good, up to 4 for the second unit, and so on). It is in general not optimal to sell all units to the agent with the higher reservation price for the first, or average unit. Instead, an optimal mechanism will ensure that units are distributed in such a way which maximises the seller's revenue (see Varian (1996), for a numerical example which illustrates the last point). We believe that such considerations are necessary in order for auction protocols to be implemented in a general purpose multi-agent system, like ADEPT.

We also provide an alternative efficient protocol for the case where the ss-agent fails, for some reason, to initiate it. The idea that agents may fail to initiate a mechanism which is beneficial for them is alien to the economic literature. However, in the context of open multi-agent environments, like the Internet, this must be considered. The pre-auction protocol which we suggest goes some way towards restoring this lost efficiency. Notice that the insurance payments, which are necessary to make sp-agents better-off participating in the pre-auction, reduce the overall efficiency of this protocol compared to an English auction. However, if these payments are small enough, the protocol still leaves everyone better off using it compared to the option of one-to-one negotiations.

Finally, it should be noted that automated auctions already take place on the Internet<sup>12</sup> and that most of them tend to use an English auction. However, this invariably has more to do with the familiarity of the protocol, than with the details of a rigorous analysis, like the one presented in this paper.

## 8 Conclusions and Future Work

We presented two examples for efficient negotiation protocols for service allocation in multi-agent systems. Our

work is set within the domain of business process management, and is specifically tailored for the requirements of the ADEPT system. Both protocols address the desiderata laid out in section 2, and the usual efficiency criteria from mechanism design. The first protocol, a multi-dimensional English auction, is shown to be a quick and efficient method for a service-seeking agent to achieve its goals, when re-sale of services is not possible. The second protocol is designed for the use of service-providing agents, who are not satisfied with the length and outcome of one-to-one negotiations. This "pre-auction" agreement is shown to be efficient under the same set of criteria. Finally, we draw some general conclusions on how negotiating agents might be designed to take into account the multi-farious results from Game Theory and mechanism design. In an open environment, it seems highly unlikely that negotiating agents will be able to calculate their best-response to any protocol they might encounter. Instead, we suggest a classification system, based on equivalence conditions from mechanism design. We describe some of the properties of such a classifying system, and show that it can also be used as the basis of decisions whether to initiate protocols.

For the future, we seek to extend our analysis of negotiation services to the more difficult one-to-one case (under the same set of constraints discussed in this paper). Second, we intend to start implementing negotiating agents, in line with the philosophy of section 6. We will then compare the relative merits of the heuristic approach currently employed in ADEPT, with the game theoretic approach advocated herein, in order to ascertain the operational strengths and weaknesses of the two approaches.

## References

- ADMATI A. R. AND PERRI M. (1987), "Strategic delay in bargaining", *Review of Economic Studies*, **54**: 345-364.
- BINMORE K. (1992), *Fun and Games*. D. C. Heath, Lexington, MA.
- BINMORE K. (1985), "Bargaining and Coalitions", in Alvin E. Roth (ed.) *Game Theoretic Models of Bargaining*, Cambridge University Press.
- BINMORE K., AND N. VULKAN, (1997), "Applying Game Theory to Automated Negotiation", *DIMACS Workshop on Economics, Game Theory and the Internet*, April 1997, Rutgers University, New Brunswick, NJ.
- JENNINGS N. R., P. FARATIN, M. J. JOHNSON, T. J. NORMAN, P. O'BRIEN AND M. E. WIEGAND (1996) "Agent-Based Business Process Management", *International Journal of Cooperative Information Systems*, **5** (2 & 3): 105-130.
- KRAUS S. (1997), "Negotiation and Cooperation in Multi-Agent Environments", *Artificial Intelligence*, **94**, 79-97.
- KREPS D. M. (1990), *Game Theory and Economic Modelling*. Oxford: Clarendon Press.
- LEE L. C. (1996), "Progressive Multi-Agent Negotiation", In *Proceedings of the Second International Conference on Multi-Agent Systems (ICMAS-96)*, 173-180, Kyoto, Japan.

<sup>12</sup>Such auctions are being used to trade a variety of goods including: vintage records (<http://www.infohwy.com/nauck/vra19/PROTOCOL.htm>), computers and electronics (<http://www.onsale.com>), art (<http://www.7cs.com>), and objects in general (<http://www.auctionline.com>; <http://www.interauction.com>; <http://auction.eecs.umich.edu>).

- McMILLAN J. (1994), "Selling spectrum rights", *Journal of Economic Perspectives*, **8**, 145-162.
- MOORE J. (1992), "Implementation, Contracts, and Renegotiation in Environments with Complete Information", pp. 182-282 in *Advances in Economic Theory*, Volume I (J.-J. Laffont, ed.), Cambridge: Cambridge University Press.
- MUELLER H. J. (1996), "Negotiation Principles", in *Foundations of Distributed Artificial Intelligence*, eds. G.M.P. O'Hare and N.R. Jennings, 211-229, Wiley.
- PEARCE D. G. (1984), "Rationalizable Strategies Behavior and the Problem of Perfection", *Econometrica*, **52**, 1029-1050.
- PRUITT D. G. (1981), *Negotiation Behaviour*, Academic Press.
- PRUITT D. G. AND P. J. CARNEVALE (1993), *Negotiation in Social Conflict*, Open University Press.
- RAIFFA H. (1982), *The Art and Science of Negotiation*, Harvard University Press.
- RODRIGUEZ, J. A., NORRIEGA P., SIERRA C., AND J. PADGET (1997), "FM96.5 A Java-based Electronic Auction House", *Proc. Second International Conference on The Practical Application of Intelligent Agents and Multi-Agent Technology (PAAM-97)*, 207-224.
- ROSENSCHEIN J. S., AND G. ZLOTKIN (1994), *Rules of Encounter*, MIT press.
- SANDHOLM T. W. AND V. R. LESSER (1995), "Issues in automated negotiation and electronic commerce: Extending the contract net framework", In *Proceedings of the First International Conference on Multi-Agent Systems (ICMAS-95)*, 328-335, San Francisco, CA.
- SANDHOLM T. W. (1996), "Limitation of the Vickrey Auction in Computational Multi-agent Systems", *Proceedings of the Second International Conference on Multi-Agent Systems (ICMAS-96)*, 299-306, Kyoto, Japan.
- SIERRA, C., P. FARATIN, AND N. R. JENNINGS (1997) "A Service-Oriented Negotiation Model between Autonomous Agents", In *Multi-Agent Rationality*, Eds. M. Boman and W. Van de Velde, LNAI-1237, Springer Verlag.
- VARIAN R. H. (1995), "Economic Mechanism Design for Computerised Agents", *USENIX Workshop on Electronic Commerce*, July 1995, New York.