Agents Contracting Tasks in Non-Collaborative Environments*

Sarit Kraus
Department of Mathematics and Computer Science
Bar Ilan University Ramat Gan, 52900 Israel
sarit@bimacs.cs.biu.ac.il

Abstract

Agents may sub-contract some of their tasks to other agent(s) even when they don't share a common goal. An agent tries to contract some of its tasks that it can't perform by itself, or when the task may be performed more efficiently or better by other agents. A "selfish" agent may convince another "selfish" agent to help it with its task, even if the agents are not assumed to be benevolent, by promises of rewards. We propose techniques that provide efficient ways to reach subcontracting in varied situations: the agents have full information about the environment and each other vs. subcontracting when the agents don't know the exact state of the world. We consider situations of repeated encounters, cases of asymmetric information, situations where the agents lack information about each other, and cases where an agent subcontracts a task to a group of agents. We also consider situations where there is a competition either among contracted agents or contracting agents. In all situations we would like the contracted agent to carry out the task efficiently without the need of close supervision by the contracting agent. The contracts that are reached are simple, Pareto-optimal and stable.

Introduction

Research in Distributed Problem Solvers assumes that it is in the agents' interest to help one another. This help can be in the form of the sharing of tasks, results, or information [Durfee, 1992]. In task sharing, an agent with a task it cannot achieve on its own will attempt to pass the task, in whole or in part, to other agents, usually on a contractual basis [Davis and Smith, 1983]. This approach assumes that agents not otherwise occupied will readily take on the task. Similarly, in information or result sharing, information is shared among agents with no expectation of a return [Lesser, 1991; Conry *et al.*, 1990]. This benevolence is based on the assumption common to many approaches to coordination: That the goal is for the system to solve the problem as best it can, and therefore the agents have a shared, often implicit, global goal that they are all unselfishly committed to achieving.

It was observed in [Grosz and Kraus, 1993] that agents may sub-contract some of their tasks to other agents also in environments where the agents do not have a common goal and there is no globally consistent knowledge.[1] That is, a selfish agent that tries to carry out its own individual plan in order to fulfill its own tasks may sub-contract some of its tasks to another selfish agent(s). An agent tries to contract some of its tasks that it can't perform by itself, or when the task may be performed more efficiently or better by other agents. The main question is how an agent may convince another agent to do something for it when the agents don't share a global task and the agents are not assumed to be benevolent. Furthermore, we would like the contracted agent to carry out the task efficiently without the need of close supervision by the contracting agent. This will enable the contracting agent to carry out other tasks simultaneously.

There are two main ways to convince another selfish agent to perform a task that is not among its own tasks: threats to interfere with the agent carrying out its own tasks or promises of rewards. In this paper we concentrate on subcontcracting by rewards. Rewards may be in two forms. In the first approach one agent may promise to help the other in its tasks in the future in return for current help. As was long ago observed in economics, barter is not an efficient basis for cooperation. In a multi-agent environment, an agent that wants to subcontract a task to another agent may not be able to help it in the future, or one agent that may be able to help in another agent's task may not need help in carrying out its own tasks. In the second approach a monetary system is developed that is used for rewards. The rewards can be used later for other purposes. We will show that a monetary system

*This material is based upon work supported by the National Science Foundation under Grant No. IRI-9123460. I would like to thank Jonathan Wilkenfeld for his comments.

[1]Systems of agents acting in environments where there is no global common goal (e.g., [Sycara, 1990; Zlotkin and Rosenschein, 1991; Kraus *et al.*, 1991; Ephrati and Rosenschein, 1991]) are called *Multi-Agent Systems* [Bond and Gasser, 1988; Gasser, 1991].

for the multi-agent environment that allows for side payments and rewards between the agents yields an efficient contracting mechanism. The monetary profits may be given to the owners of the automated agents. The agents will be built to maximize expected utilities that increase with the monetary values, as will be explained below.

The issue of contracts has been investigated in economics and game-theory in the last two decades (e.g., [Arrow, 1985; Ross, 1973; Rasmusen, 1989; Grossman and Hart, 1983; Hirshleifer and Riley, 1992]). They have considered situations in which a person or a company contracts a task to another person or company. In this paper we adjust the models that were developed in economics and game-theory to fit distributed artificial intelligence situations.

We will consider varied situations: In the Section *Contracts Under Certainty* one agent subcon tracts a task to another one when the agents have full information about the environment and each other. In the Section *Contracts Under Uncertainty*, we consider contracting when the agents don't know the exact state of the world. The situation in which an agent may subcontract its tasks several times to the same agent is considered in Section *Repeated Encounters*, and situations of asymmetric information or when the agents lack information about each other is dealt with in the Section *Asymmetric and Incomplete Information*. We conclude with the case of an agent subcontracting a task to a group of agents. In all these cases, we consider situations where the contracting agent doesn't supervise the contracted agents' performance and situations where there is a competition among possible contracted agents or possible contracting agents.

Preliminaries

We will refer to the agent that subcontracts one of its tasks to another agent as the *contracting agent* and to the agent that agrees to carry out the task as the *contracted agent*. The *effort* level is the time and work intensity which the contracted agent puts into fulfilling the task. We denote the set of all possible efforts by E. In all cases, the contracted agent chooses how much effort to extend, but its decision may be influenced by the contract offered by the contracting agent. We assume that there is a monetary value $q(e)$ of performing a task which increases with the effort involved. That is, the more time and effort put in by the contracted agent, the better the outcome. The contracting agent will pay the contracted agent a wage w (which can be a function of q). There are several properties we require from our mechanism for subcontracting:

Simplicity: The contract should be simple and there should be an algorithm to compute it.

Pareto-Optimality: There is no other contracted arrangement that is preferred by both sides over the one they have reached.

Stability: We would like the results to be in equilibrium and that the contracts will be reached and executed without delay.

Concerning the simplicity and stability issues, there are two approaches for finding equilibria in the type of situations under consideration here [Rasmusen, 1989]. One is the straight game theory approach: a search for Nash strategies or for perfect equilibrium strategies. The other is the economist's standard approach: set up a maximization problem and solve using calculus. The drawback of the game theory approach is that it is not mechanical. Therefore, in our previous work on negotiation under time constraints, we have identified perfect-equilibrium strategies and proposed to develop a library of meta strategies to be used when appropriate [Kraus and Wilkenfeld, 1991a; Kraus and Wilkenfeld, 1991b]. The maximization approach is much easier to implement. The problem with the maximization approach in our context is that players must solve their optimization problems jointly: the contracted agent's strategy affects the contracting agent's maximization problem and vice versa. In this paper we will use the maximization approach, with some care, by embedding the contracted agent's maximization problem into the contracting agent's problem as a constraint. This maximization problem can be solved automatically by the agent.

The agents' utility functions play an important role in finding an efficient contract. As explained above, we propose to include a monetary system in the multi-agent environment. This system will provide a way for providing rewards. However, it is not always the case that the effort of an agent can be assigned the same monetary values. Each designer of an automated agent needs to provide its agent with a decision mechanism based on some given set of preferences. Numeric representations of these preferences offer distinct advantages in compactness and analytic manipulation [Wellman and Doyle, 1992]. Therefore, we propose that each designer of autonomous agents will develop a numerical utility function that it would like its agent to maximize. In our case the utility function will depend on monetary gain and effort. This is especially important in situations where there is uncertainty in the situation and the agents need to make decisions under risk considerations. Decision theory offers a formalism for capturing risk attitudes. If an agent's utility function is concave, it is risk averse. If the function is convex, it is risk prone, and a linear utility function yields risk neutral behavior [Hirshleifer and Riley, 1992].

We denote the contracted agent's utility function by U which is a decreasing function in effort and an increasing function in wage w. We assume that if the contracted agent won't accept the contract from the contracting agent, its utility, (i.e., its reservation price) which is known to both agents is $\hat{u}$. This outcome can result either from not doing anything or performing some other tasks at the same time. We denote the contracting agent's utility function by V and it is an increasing function with the value of performing the task (q) and decreasing function with the wage w paid to the contracted agent. In our system we assume that the contracting agent rewards the contracted agent *after* the task is carried out. In such situations there should be a technique for enforcing this reward. In case of multiple encounters reputational considerations may yield appropriate behavior. In a single encounter some external intervention may be required to enforce commitments.

Contracts Under Certainty

In this case we assume that all the relevant information about the environment and the situation is known to both agents. In the simplest case the contracting agent can observe and supervise the contracted agent's effort and actions and force it to make the effort level preferred by the contracting agent by paying it only in case it makes the required effort. The amount of effort required from the contracted agent will be the one that maximizes the contracting agent's outcome, taking into account the task fulfillment and the payments it needs to make to the contracted agent.

However, in most situations it is either not possible or too costly for the contracting agent to supervise the contracted agent's actions and observe its level of effort. In some cases, it may be trying to carry out another task at the same time, or it can't reach the site of the action (and that is indeed the reason for subcontracting). If the outcome is a function of the contracted agent's effort and if this function is known to both agents the contracting agent can offer the contracted agent a *forcing contract* [Harris and Raviv, 1978; Rasmusen, 1989]. In this contract, the contracting agent will pay the contracted agent only if it provides the outcome required by the contracting agent. If the contracted agent accepts the contract, he will perform the task with the effort that the contracting agent finds to be most profitable to itself even without supervision. Note that the outcome won't necessarily be with the highest effort on the part of the contracted agent, but rather the effort which provides the contracting agent with the highest outcome. That is, the contracting agent should pick an effort level e^* that will generates the efficient output level q^*. Since we assume that there are several possible agents available for contracting, in equilibrium, the contract must provide the contracted agent with the utility $\hat{u}$.[2] The contracting agent needs to choose a wage function such that $U(e^*, w(q^*)) = \hat{u}$ and $U(e, w(q)) < \hat{u}$ for $e \neq e^*$. We demonstrate this case in the following example.

Example 1: Contracting Under Certainty

The US and Germany have sent several mobile robots independently to Mars to collect minerals and ground samples and to conduct experiments. One of the US robots has to dig some minerals on Mars far from the other US robots. There are several German robots in that area and the US robot would like to subcontract some of its digging. The US robot approaches one of the German robots that can dig in three levels of effort (e): Low, Medium and High denoted by 1,2 and 3 respectively. The US agent can't supervise the German robot's effort since it wants to carry out another task simultaneously. The value of digging is $q(e) = \sqrt{100e}$. *The US robot's utility function, if a contract is reached, is* $V(q, w) = q - w$ *and the German robot's utility function in case it accepts the contract is* $U(e, w) = 17 - \frac{10}{w} - 2e$, *where* w *is the payment to the German robot. If the German robot rejects the contract, it will busy itself with maintenance tasks and its utility will be 10. It is easy to calculate that the best effort level from the US robot's point of view is 2, in which there will be an outcome of* $\sqrt{200}$. *The contract that the US robot offers to the German robot is* $3\frac{1}{3}$ *if the outcome is* $\sqrt{200}$ *and 0 otherwise. This contract will be accepted by the German robot and its effort level will be Medium.*

There are two additional issues of concern. The first one is how the contracting agent chooses which agent to approach. In the situation of complete information (we consider the incomplete information case in Section *Asymmetric and Incomplete Information*) it should compute the expected utility for itself from each contract with each agent and chooses the one with the maximal expected utility.

Our model is also appropriate in the case in which there are several contracting agents, but only one possible contracted agent. In such a case, there should be information about the utilities of the contracting agents in the event that they don't sign a contract. The contracted agent should compute the level of effort that maximizes its expected utility (similar to the computation of the contracting agent in the reverse case) and make an offer to the contracting agent that will maximize its outcome.

Contracts Under Uncertainty

In most subcontracting situations, there is uncertainty concerning the outcome of an action. If the contracted agent chooses some effort level, there are several possibilities for an outcome. For example, suppose an agent on Mars subcontracts digging for samples of a given mineral and suppose that there is an uncertainty about the depth of the given mineral at the site. If the contracted agent chooses a high effort level but the mineral level is deep underground the outcome may be similar to the case where the contracted agent chooses a low level of effort but the mineral is located near the surface. But, if it chooses a high effort level when the mineral is located near the surface, the outcome may be much better. In such situations the outcome of performing a task doesn't reveal the exact effort level of

[2] We assume that if the contracted agent is indifferent between two actions, it will choose the one preferred by the contracting agent.

the contracted agent and choosing a stable and maximal contract is much more difficult.

We will assume that the world may be in one of several states. Neither the contracting agent nor the contracted agent knows the exact state of the world when agreeing on the contract, as well as when the contracted agent chooses the level of effort to take, after agreeing on the contract. The contracted agent may observe the state of the world *after* choosing the effort level (during or after completing the task), but the contracting agent can't observe it. For simplicity, we also assume that there is a set of possible outcomes to the contracted agent carrying out the task $Q = \{q_1, ..., q_n\}$ such that $q_1 < q_2 < ... < q_n$ that depends on the state of the world and the effort level of the contracted agent. Furthermore, we assume that given a level of effort, there is a probability distribution that is attached to the outcomes that is known to both agents.[3] Formally, we assume that there is a probability function $\wp : E \times Q \rightarrow I\!R$, such that for any $e \in E$, $\sum_1^n \wp(e, q_i) = 1$ and for all $q_i \in Q$, $\wp(e, q_i) > 0$.[4] The contracting agent's problem is to find a contract that will maximize its expected utility, knowing that the contracted agent may reject the contract or if it accepts the contract the effort level is chosen later [Rasmusen, 1989]. The contracting agent's payment to the contracted agent can be based only on the outcome. Let us assume that in the contract that will be offered by the contracting agent, for any q_i $i = 1, ..., n$ the contracting agent will pay the contracted agent w_i. The maximization problem can be constructed as follows (see also [Rasmusen, 1989]).

$$Maximize_{w_1,...w_n} \sum_1^n \wp(\hat{e}, q_i) V(q_i, w_i) \qquad (1)$$

with the constraints:

$$\hat{e} = argmax_{e \in E} \sum_1^n \wp(e, q_i) U(e, w_i) \qquad (2)$$

$$\sum_1^n \wp(\hat{e}, q_i) U(\hat{e}, w_i) > \hat{u} \qquad (3)$$

Equation (1) states that the contracting agent tries to choose the payment to the contracted agent so as to maximize its expected utility subject to the constraint that the contracted agent will prefer the contract over rejecting it (3) and that the contracted agent prefers the effort level that the contracting agent prefers, given the contract it is offered (2).

The main question is whether there is an algorithm to solve this maximization problem and whether such a contract exists. This depends primarily on the utility functions of the agents. If the contracting agent and the contracted agent are risk neutral, then solving the maximization problem can be done using any linear programming technique (e.g, simplex, see for example [Pfaffenberger and Walker, 1976].) Furthermore, in most situations, the solution will be very simple: the contracting agent will receive a fixed amount of the outcome and the rest will go to the contracted agent. That is, $w_i = q_i - C$ for $1 \leq i \leq n$, where the constant C is determined by constraint (3) [Shavell, 1979].

Example 2: Risk Neutral Agents Under Uncertainty

Suppose the utility function of the German robot from Example 1 is $U(w, e) = w - e$ and that it can choose between two effort levels Low (e=1) and High (e=2) and its reservation price is $\hat{u} = 1$. There are two possible monetary outcomes to the digging: $q_1 = 8$ and $q_2 = 10$ and the US robot's utility function is as in the previous example, i.e., $V(q, w) = q - w$.

If the German robot chooses the Lower level effort then the outcome will be q_1 with probability $\frac{3}{4}$ and q_2 with probability $\frac{1}{4}$ and if it takes the High level effort the probability of q_1 is $\frac{1}{8}$ and of q_2 it is $\frac{7}{8}$. In such situations, the US robot should reserve to itself a profit of $6\frac{3}{4}$. That is, $w_1 = 1\frac{1}{4}$ and $w_2 = 3\frac{1}{4}$. The German robot will choose the High level effort.

If the agents are not neutral toward risk, the problem is much more difficult. However, if the utility function for the agents are carefully chosen, an algorithm does exist. Suppose the contracted agent is risk averse and the contracting agent is risk neutral (the methods are also applicable when both are risk averse). Grossman and Hart [Grossman and Hart, 1983] presented a three-step procedure to find appropriate contracts. The first step of the procedure is to find for each possible effort level the set of wage contracts that induce the contracted agent to choose that effort level. The second step is to find the contract which supports that effort level at the lowest cost to the contracting agent. The third step is to choose the effort level that maximizes profits, given the necessity to support that effort with a costly wage contract. For space reasons, we won't present the formal details of the algorithm here, and also in the rest of the paper.

Repeated Encounters

Suppose the contracting agent wants to subcontract its tasks several (finite) times. Repetition of the encounters between the contracting and the contracted agents enables the agents to reach efficient contracts if the

[3] A practical question is how the agents find the probability distribution. It may be that they have preliminary information about the world, e.g., what is the possibility that a given mineral will be in that area of Mars. In the worst case, they may assume an equal distribution. The model can be easily extended to the case that each agent has different beliefs about the state of the world [Page, 1987].

[4] The formal model in which the outcome is a function of the state of the world and the contracted agent's effort level, and in which the probabilistic function gives the probability of the state of the world which is independent of the contracted agent's effort level is a special case of the model described here [Page, 1987; Ross, 1973; Harris and Raviv, 1978].

umber of encounters is large enough. The contracting agent could form an accurate estimate of the contracted agent's effort, based on the average outcome, over time. That is, if the contracting agent wants the contracted agent to take a certain effort level $\hat{e}$, in all the encounters, it can compute the expected outcome over time if the contracted agent actually performs the task with that effort level. The contracting agent can keep track of the cumulative sum of the actual outcomes and compare it with the expected outcome. If there is some time T in which the outcome is below a given function of the expected outcome, the contracting agent should impose a big "punishment" on the contracted agent. If the function over the expected outcome is chosen carefully [Radner, 1981], the probability of imposing a "punishment" when the contracted agent is in fact carrying out the desired effort level can be made very low, while the probability of eventually imposing the "punishment" if the agent doesn't do e is one.

Asymmetric and Incomplete Information

In some situations the contracting agent does not know the utility function of the contracted agent. The contracted agent may be one of several types that reflect the contracted agent's ability to carry out the task, its efficiency or the cost of its effort. However, we assume that given the contracted agent's type, its utility function is known to its opponent. For example, suppose Germany builds robots of two types. The specifications of the robots are known to the German robots and to the US robots, however, the US robots don't know the specific type of the German robots that they meet.

As in previous sections the output is a function of the contracted agent's effort level, and the probability function $\wp$ indicates the probability of each outcome, given the effort level and the agent's type. The question remains which contract the contracting agent should offer when it doesn't know the contracted agent's type. A useful technique in such situations is for the contracting agent to search for an optimal mechanism [Demougin, 1989] as follows: the contracting agent offers the contracted agent a menu of contracts that are functions of its type and the outcome. The agents chooses a contract (if at all) and announces it to the contracting agent. Given this contract, the contracted agent chooses an effort level which maximizes its own expected utility. In each of the menu's contracts, the contracted agent's expected utility should be at least as its expected utility if it doesn't sign the contract. We also concentrate only on contracts in which it will always be in the interest of the contracted agent to honestly report its type. It was proven that this requirement is without loss of generality [Myerson, 1982]. It was also shown that in some situations, an efficient contract can be reached without communication [Demougin, 1989], but we omit the discussion here for space reasons.

If there are several agents whose types are unknown to the contracting agent and it must choose among them, the following mechanism is appropriate: The contracting agent announces a set of contracts based on the agent's type and asks the potential contracted agents to report their types. On the basis of these reports the contracting agent chooses one agent [McAfee and McMillan, 1987].

In other situations, the contracting agent knows the utility function of the contracted agent, but the contracted agent is able to find more information on the environment than the contracting agent. For example, when the German robot reaches the area where it needs to dig, it determines the structure of this area. This information is known only to the German robot and not to the US robot. The mechanism that should be used in this context is the following: The contracting agent offers a payment arrangement which is based on the outcome and the message the contracted agent will send to the contracting agent about the additional information it possesses. If the contracted agent accepts the offer, it will observe the information (by going to the area, or using any of its sensors etc.). Then it will send a message to the contracting agent and will choose its effort level. Eventually, after the task is finished and the outcome is observed, the contracting agent will pay the rewards. Also in this case [Christensen, 1981], the agents can concentrate on the class of contracts that induce the contracted agent to send a truthful message. This is since for any untruthful contracts, a truthful one can be found in which the expected utility of agents is the same. A maximization solvable problem can be constructed here, but we omit it for space reasons.

Subcontracting to a Group

Suppose that the task the contracting agent wants to contract for can be performed by a group of agents. Each of the contracted agents is independent in the sense that it tries to maximize its own utility. The contracting agent offers a contract to each of the possible contracted agents. If one of them rejects the offer, then the contracting agent cannot subcontract the task. Otherwise, the contracted agents simultaneously choose effort levels.

In other situations, the contracting agent can't observe the individual outcome (or such an outcome does not exists) but rather observe only the overall outcome from the effort of all agents [Holmstrom, 1982].Here, even in the case of certainty, i.e., the state of the world is known, there is a problem in making the contracted agents take the preferred level of action, since there is no way for the contracting agent to find out the effort level of each of the individual agent, given the overall output. For example, suppose two robots agreed to dig minerals, but they both put the minerals in the same truck, so it is not possible to figure out who digs what.

If the contracting agent wants the contracted agents to take the vector of levels effort e^* it can search for a contract such that, if the outcome is $q \geq q(e^*)$ then $w_i(q) = b_i$ and otherwise 0, such that $U(e_i^*, b_i) \geq \hat{u}_i$. That is, if all agents choose the appropriate effort level, each of them gets b_i and if any of them does not, all get nothing.

Conclusions

In this paper we presented techniques that can be used in different cases where sub-contracting of a task by an agent to another agent or a set of agents in non-collaborative environments is beneficial. In all the situations, simple Pareto-optimal contracts can be found by using techniques of maximization with constraints. In the case where the agents have complete information about each other, there is no need for negotiations and a contract is reached without a delay even when the contracting agent doesn't supervise the contracted agent's actions. If there is asymmetric information, or the agents are not sure about their opponents' utility functions, a stage of message exchange is needed to reach a contract.

References

Arrow, K. J. 1985. The economics of agency. In Pratt, J. and Zeckhauser, R., editors 1985, *Principals and Agents: The Structure of Business*. Harvard Business School Press. 37–51.

Bond, A. H. and Gasser, L. 1988. An analysis of problems and research in DAI. In Bond, A. H. and Gasser, L., editors 1988, *Readings in DAI*. Morgan Kaufmann Pub., Inc., Ca. 3–35.

Christensen, J. 1981. Communication in agencies. *Bell Journal of Economics* 12:661—674.

Conry, S.E.; MacIntosh, D. J.; and Meyer, R.A. 1990. DARES: A distributed automated REasoning system. In *Proc. of AAAI90*, MA. 78–85.

Davis, R. and Smith, R.G. 1983. Negotiation as a metaphor for distributed problem solving. *Artificial Intelligence* 20:63–109.

Demougin, D. 1989. A renegotiation-proof mechanism for a principle-agent model with moral hazard and adverse selection. *The Rand Journal of Economics* 20:256–267.

Durfee, E. 1992. What your computer really needs to know, you learned in kindergarten. In *Proc. of AAAI-92*, California. 858–864.

Ephrati, E. and Rosenschein, J. 1991. The clarke tax as a consensus mechanism among automated agents. In *Proc. of AAAI-91*, California. 173–178.

Gasser, L. 1991. Social concepts of knowledge and action: DAI foundations and open systems semantics. *Artificial Intelligence* 47(1–3):107–138.

Grossman, S. and Hart, O. 1983. An anaysis of the principal-agent problem. *Econometrica* 51(1):7–45.

Grosz, B. and Kraus, S. 1993. Collaborative plans for group activities. In *IJCAI93*, French.

Harris, M. and Raviv, A. 1978. Some results on incentive contracts with applications to education and employment, health insurance, and law enforcement. *The American Economic Review* 68(1):20–30.

Hirshleifer, J. and Riley, J. 1992. *The Analytics of Uncertainty and Information*. Cambridge University Press, Cambridge.

Holmstrom, B. 1982. Moral hazard in teams. *Bell Journal of Economics* 13(2):324—340.

Kraus, S. and Wilkenfeld, J. 1991a. The function of time in cooperative negotiations. In *Proc. of AAAI-91*, California. 179–184.

Kraus, S. and Wilkenfeld, J. 1991b. Negotiations over time in a multi agent environment: Preliminary report. In *Proc. of IJCAI-91*, Australia. 56–61.

Kraus, S.; Ephrati, E.; and Lehmann, D. 1991. Negotiation in a non-cooperative environment. *J. of Experimental and Theoretical AI* 3(4):255–282.

Lesser, V.R. 1991. A retrospective view of fa/c distributed problem solving. *IEEE Transactions on Systems, Man, and Cybernetics* 21(6):1347–1362.

McAfee, R. P. and McMillan, J. 1987. Competition for agency contracts. *The Rand Journal of Economics* 18(2):296–307.

Myerson, R. 1982. Optimal coordination mechanisms in generalized principal-agent problem. *Journal of Mathematical Economics* 10:67–81.

Page, F. 1987. The existence of optimal contracts in the principal-agent model. *Journal of Mathematical Economics* 16(2):157–167.

Pfaffenberger, R. and Walker, D. 1976. *Mathematical Programming for Economics and Business*. The IOWA State University Press, Ames, IOWA.

Radner, R. 1981. Monitoring cooperative agreements in a repeated principal-agent relationships. *Econometrica* 49(5):1127–1148.

Rasmusen, E. 1989. *Games and Information*. Basil Blackwell Ltd., Cambridge, Ma.

Ross, S. 1973. The economic theory of agency: The principal's problem. *The American Economic Review* 63(2):134–139.

Shavell, S. 1979. Risk sharing and incentives in the principal and agent relationship. *Bell Journal of Economics* 10:55–79.

Sycara, K. P. 1990. Persuasive argumentation in negotiation. *Theory and Decision* 28:203–242.

Wellman, M. and Doyle, J. 1992. Modular utility representation for decision-theoretic planning. In *Proc. of AI planning Systems*, Maryland. 236—242.

Zlotkin, G. and Rosenschein, J. 1991. Incomplete information and deception in multi-agent negotiation. In *Proc. IJCAI-91*, Australia. 225–231.